NORTH ATLANTIC TREATY ORGANISATION

RESEARCH AND TECHNOLOGY ORGANISATION

AC/323(IST-021)TP/43

www.rto.nato.int

RTO TECHNICAL REPORT **TR-IST-021**

Multimedia Visualisation of Massive Military Datasets

(Visualisation multimédia des jeux de données militaires massifs)

This Report documents the findings of Task Group IST-021/RTG-007.

Published June 2007

Distribution and Availability on Back Cover

NORTH ATLANTIC TREATY ORGANISATION

RESEARCH AND TECHNOLOGY ORGANISATION

AC/323(IST-021)TP/43

www.rto.nato.int

RTO TECHNICAL REPORT **TR-IST-021**

Multimedia Visualisation of Massive Military Datasets

(Visualisation multimédia des jeux de données militaires massifs)

This Report documents the findings of

Task Group IST-021/RTG-007.

The Research and Technology Organisation (RTO) of NATO

RTO is the single focus in NATO for Defence Research and Technology activities. Its mission is to conduct and promote co-operative research and information exchange. The objective is to support the development and effective use of national defence research and technology and to meet the military needs of the Alliance, to maintain a technological lead, and to provide advice to NATO and national decision makers. The RTO performs its mission with the support of an extensive network of national experts. It also ensures effective co-ordination with other NATO bodies involved in R&T activities.

RTO reports both to the Military Committee of NATO and to the Conference of National Armament Directors. It comprises a Research and Technology Board (RTB) as the highest level of national representation and the Research and Technology Agency (RTA), a dedicated staff with its headquarters in Neuilly, near Paris, France. In order to facilitate contacts with the military users and other NATO activities, a small part of the RTA staff is located in NATO Headquarters in Brussels. The Brussels staff also co-ordinates RTO's co-operation with nations in Middle and Eastern Europe, to which RTO attaches particular importance especially as working together in the field of research is one of the more promising areas of co-operation.

The total spectrum of R&T activities is covered by the following 7 bodies:

- AVT Applied Vehicle Technology Panel
- HFM Human Factors and Medicine Panel
- IST Information Systems Technology Panel
- NMSG NATO Modelling and Simulation Group
- SAS System Analysis and Studies Panel
- SCI Systems Concepts and Integration Panel
- SET Sensors and Electronics Technology Panel

These bodies are made up of national representatives as well as generally recognised 'world class' scientists. They also provide a communication link to military users and other NATO bodies. RTO's scientific and technological work is carried out by Technical Teams, created for specific activities and with a specific duration. Such Technical Teams can organise workshops, symposia, field trials, lecture series and training courses. An important function of these Technical Teams is to ensure the continuity of the expert networks.

RTO builds upon earlier co-operation in defence research and technology as set-up under the Advisory Group for Aerospace Research and Development (AGARD) and the Defence Research Group (DRG). AGARD and the DRG share common roots in that they were both established at the initiative of Dr Theodore von Kármán, a leading aerospace scientist, who early on recognised the importance of scientific support for the Allied Armed Forces. RTO is capitalising on these common roots in order to provide the Alliance and the NATO nations with a strong scientific and technological basis that will guarantee a solid base for the future.

The content of this publication has been reproduced
directly from material supplied by RTO or the authors.

Published June 2007

ISBN 978-92-837-0067-8

Table of Contents

List of Figures

List of Tables

IST-021 Membership List

CANADA

Dr. Justin G. HOLLANDS
Defence R&D Canada Toronto
1133 Sheppard Avenue West
P.O. Box 2000
Toronto, Ontario M3M 3B9

Dr. Zachary JACOBSON
Health Canada Applied Research and Analysis Directorate
Jeanne Mance Building, Floor 5
Tunney's Pasture, Postal Locator 1905C
Ottawa, Ontario K1A 0K9

Ms. Sarah ROSSER
Health Canada
2 793 1/2 Bank Street
Ottawa, Ontario K1S 3V7

Dr. Maurice Martin TAYLOR
Martin Taylor Consulting
369 Castlefield Avenue
Toronto, Ontario M5N 1L4

Mr. Vincent K. TAYLOR
Defence R&D Canada Ottawa
Network Information Operations
3701 Carling Avenue
Ottawa, Ontario K1A 0Z4

Dr. Malcolm R. VANT
Ottawa Director General
Defence R&D Canada Ottawa
3701 Carling Avenue
Ottawa, Ontario K1A 0Z4

CZECH REPUBLIC

Dr. Prof. Ladislav NESRSTA
DELINFO
Chodska 15
612 00 BRNO

DENMARK

Mrs. Lisbeth M. RASMUSSEN
Danish Defence Acquisition and Logistics Organization (DALO)
FMT-C-TDBE12
Lautrupbjerg 1
DK 2750 Ballerup

FRANCE

ICT Eric PFANNSTIEL
DGA/SPOTI
Fort d'Issy-les-Moulineaux
18 rue du Docteur Zamenhof
92131 Issy-les-Moulineaux

GEORGIA

Mrs. Jr. Lt. Nino KHATISKATSI
Central Board of Strategy Planning and Scientific Technical Researches
Moscow avenue, 7/1
380020 Tbilissi

GERMANY

Dipl-Ing Annette KASTER
Forschungsgesellschaft für Angewandte (FGAN)
Neuenahrerstr. 20
D-53343 Wachtberg-Werthhoven

NORWAY

Mr. Tore SMESTAD
Norwegian Defence Research Establishment (FFI)
PO Box 25
NO-2027 Kjeller

PORTUGAL

Maj. José Carlos DA SILVA VERISSIMO
PO Ministry of Defence
Direccao do Servico de Transmissoes
Avenida Infante Santo, 49, 5o andar
1350-177 Lisboa

SLOVAK REPUBLIC

Lt-Col. Ir. Dusan CABALA
Ministry of Defence
Military Academy
Kt-302, PO Box 45
03101 Liptovsky Mikulas

Maj. Eng. Pavel SIMON
Military Technical Institute
CIS Department
ul. kpt. Nalepku
03101 Liptovsky Mikulas

UNITED KINGDOM

Dr. Margaret J. VARGA
QinetiQ
Spectrum Solutions Division
Malvern Technology Park
St. Andrew's Road
Malvern, Worcestershire WR14 3PS

UNITED STATES

Dr. Susan F. CHIPMAN
Office of Naval Research
Code 342
800 N. Quincy Street
Arlington, VA 22217-5660

Multimedia Visualisation of Massive Military Datasets

(RTO-TR-IST-021)

Executive Summary

The objective of IST-021/RTG-007 was to gain visualisation knowledge in common which will help facilitate the effective development of internationally useful systems. The RTG undertook four major tasks: evaluate and update visualisation systems principles developed by IST-013/RTG-002; deliver a workshop in 2002 on multimedia information visualisation, the results of which helped to guide the subsequent work of the group; propose a lecture series on Multimedia Information Visualisation for delivery in 2005 (this task was cancelled); deliver a workshop in 2004 on Visualisation and the Common Operational Picture.

Evaluations of two visualisation systems were attempted, the Norwegian NORCCIS II and the Canadian VITA. Two approaches that had been proposed to IST-021/RTG-007 (The VisTG Reference Model, and the Smestad guidelines based on information-theoretic principles) were used with NORCCIS II, but both evaluations proved complicated and were not completed in the limited time made available. Nevertheless, both attempts did find areas of potential improvement, but these improvements had already been noted by the operators and were not newly discovered by the IST-021/RTG-007 approaches. As for VITA, an "appreciation" was performed using the VisTG Reference Model that showed issues and benefits both with the evaluation process and with VITA.

Both approaches to evaluation suggested that the complexity of real system structures would usually present problems in keeping track of cross influences. Too much data was generated and too many relationships observed for a manual approach to be effective. Nevertheless:

- The VisTG Reference Model is useful as it stands as a way of thinking and talking about systems, in the same way that the ISO seven level reference model is useful in talking about communication networks.
- Using the model for a practical evaluation would require the development of automated support mechanisms.
- Understanding the data generated and relationships observed may be helped through the program of work of IST-059/RTG-025 "Visualisation Technology for Network Analysis".

Network representation and analysis appeared often as outstanding issues during the life of RTG-007. The 2002 Halden workshop, the 2003 Network of Experts workshop and the 2004 Toronto workshop deliberations and recommendations reinforced this view.

Networks are of two different kinds: "Structural" networks in which the links are physical and can be observed (e.g. information and infrastructural networks), and "Logical" networks in which the links are largely conceptual, existing in a person's mind (e.g. social networks which show the organizational relationships among their elements, or the relationship among infrastructures with respect to the effect that damage to one may have on another). Networks of relationships include causal or probabilistic networks that affect military planning. Networks are often created in a commander's mind through experience and interpretation of his displays. Such mental networks are not often available to others.

Understanding how to display these different kinds of networks effectively would help promote a common understanding within a command and improve the robustness of operational planning. IST-059/RTG-025 had been formed to study this problem.

Visualisation multimédia des jeux de données militaires massifs

(RTO-TR-IST-021)

Synthèse

L'objectif du groupe de travail IST-021/RTG-007 était d'acquérir des connaissances de visualisation en commun afin de faciliter le développement efficace de systèmes utiles à l'échelle internationale. Le RTG était chargé de quatre tâches majeures : évaluer et mettre à jour les principes des systèmes de visualisation développés par le groupe de travail IST-013/RTG-002 ; organiser un atelier en 2002 sur la visualisation d'informations multimédia dont les résultats ont contribué au travail ultérieur du groupe ; proposer une série de conférences sur la Visualisation des informations multimédia pour livraison en 2005 (cette tâche a été annulée) ; organiser un atelier en 2004 sur la Visualisation et le modèle commun opérationnel.

Deux systèmes de visualisation ont fait l'objet d'une évaluation : le norvégien NORCCIS II le canadien VITA. Les deux approches qui avaient été proposées au groupe de travail IST-021/RTG-007 (le modèle de référence VisTG et les lignes directrices Smestad s'appuyant sur les principes théoriques d'information) ont été utilisées avec NORCCIS II, mais les deux évaluations ont présenté des complications et n'ont pu être menées à bien dans le temps imparti. Des domaines d'amélioration potentielle ont toutefois été décelés par ces deux tentatives, mais ces améliorations avaient déjà été notées par les opérateurs et les approches du groupe de travail IST-021/RTG-007 n'en ont découvert aucune autre. Quant à VITA, une "estimation", réalisée à l'aide du modèle de référence VisTG, a présenté des inconvénients et des avantages à la fois avec le processus d'évaluation et VITA.

Ces deux approches d'évaluation ont indiqué que la complexité de structures de systèmes réels présente généralement des problèmes si la trace des influences croisées est conservée. De trop nombreuses données ont été générées, et relations observées, pour qu'une approche manuelle puisse être efficace. Toutefois :

- Le modèle de référence VisTG est utile dans la mesure où il constitue un moyen de penser et de parler de systèmes, de la même façon que le modèle de référence à sept niveaux ISO sert à parler de réseaux de communication.
- L'utilisation du modèle pour une évaluation pratique nécessiterait le développement de mécanismes de soutien automatisés.
- Le programme de travail du groupe IST-059/RTG-025 « Technologie de visualisation pour analyse réseau » peut permettre de mieux comprendre les données générées et relations observées.

Pendant toute la durée de vie du groupe de travail RTG-007, la représentation et l'analyse réseau sont souvent apparues comme des questions restant à régler. Les délibérations et recommandations de l'atelier à Halden en 2002, de l'atelier Réseau d'Experts en 2003 et de l'atelier de Toronto en 2004 ont renforcé cette image.

Il existe deux types de réseaux : des réseaux "structurels" dans lesquels les liaisons sont physiques et peuvent être observées (ex. : réseaux infrastructurels et d'informations), et des réseaux "logiques" dans lesquels les liaisons sont largement conceptuelles, existant dans l'esprit d'un individu (ex. : réseaux sociaux indiquant les relations organisationnelles au sein de leurs éléments, ou les relations parmi les infrastructures par rapport à l'effet que les dommages de l'un pourraient avoir sur l'autre). Les réseaux relationnels incluent des réseaux causaux ou probabilistes affectant la planification militaire. Les réseaux sont souvent créés dans un esprit de commandement, par l'expérience et l'interprétation de ses structures. Ces réseaux mentaux ne sont pas toujours accessibles aux autres.

Savoir structurer ces différents types de réseaux efficacement permettrait de tendre vers une compréhension commune à l'intérieur d'un commandement et d'améliorer la robustesse de la planification opérationnelle. Le groupe de travail IST-059/RTG-025 avait été formé pour étudier ce problème.

Chapter 1 – INTRODUCTION

V.K. Taylor
Defence Research and Development Canada

Z. Jacobson
Health Canada

M.M. Taylor
Martin Taylor Consulting, Canada

1.1 BACKGROUND AND INCEPTION

Understanding and producing information from huge amounts of rapidly changing data is a key concern in military and defence operations. The Information Technology Panel under the old NATO Defence Research Group (DRG Panel 11) envisioned that flexible and intuitive visual interfaces, possibly within Virtual Environments, could contribute greatly to helping personnel by improving both their interaction with the dataflood and their capacity to extract and manage information from it.

In 1993, Panel 11 created an Exploratory Group to determine whether automated visual information processing was an important domain for cooperative or collaborative investigation and if so, to elicit national interest in the formation of a NATO Research Study Group in that area. In May 1994, the exploratory group reported to the Panel, indicating that progress was rapid, but unfocused on military and defence needs and concerns and recommended that the Panel authorize a Workshop to bring visualisation issues more clearly to the international defence community. The Workshop, titled “Visualising non-Visual Information”, was held in November 1994.

The workshop recommended that Panel 11 form a Research Study Group to coordinate much needed research on dataflood issues related to military and defence needs, and to maintain a broad overview and a detailed view of visual information management technologies. Panel 11 agreed to the recommendation and a Research Study Group (RSG) was created in 1996 with the aim of developing methods for presenting to human users the implications of the contents of large, complex and varying military-relevant datasets of diverse kinds. The RSG would operate by meeting semi-annually in the different member nations. Shortly thereafter, Panel 11 was decommissioned and the responsibility for the RSG was transferred to DRG Panel 8, the Human Factors Panel, as its RSG-30.

RSG-30 identified seven main classes of activity in which visualisation can play an effective role: monitoring; alerting; searching; exploring; analysis; problem solving; and briefing. In each of these areas, specific and different types of interaction with information are required, thus implying that specific, different types of visualisation may be needed. The scope of RSG-30 was bounded by the aforementioned activity classes.

RSG-30 decided, from the expressions of various national interests, that “visualisation” implied understanding data as information rather than simply displaying data on a screen. Other sensors besides the eyes are frequently valuable in creating a mental picture. Hence, although derived from an initial interest in discovering and displaying the content of massive textual datasets, the horizon of RSG-30 was expanded to include non-textual material. RSG-30 interpreted visualisation as a human activity, supported by technology, by which humans make sense of complex data. The Group considered visualisation technologies, such as search engines, algorithmic processes and display devices and techniques, but only in relation to how they help humans to perform their tasks effectively. The Group emphasized the human use of the computational subsystem in ensuring that the right information be available in the form and at the time needed. In this, the Group's approach to the nature of visualisation coincides with that of the US Army, which characterizes battlefield visualisation as "*the process whereby the commander develops a clear understanding of his current state with relation to the enemy and the environment, envisions a desired end state, and then subsequently visualises the sequence of activity that will move his force from its current state to the end state*".

In 1996, RSG-30 formed an informal technical group known as the "Visualisation Network of Experts", or "Vis-NX" composed of known visualisation experts from the NATO countries. The group was to have an independent existence, but be cognizant of, and responsive to, the needs of the RSG. Vis-NX was expected to conduct a workshop annually in conjunction with a regularly scheduled meeting of the RSG and be a sounding board and advisor to the RSG on developments in visualisation science and technology.

In 1998, while RSG-30 was still active, NATO's defence science and technology structure was reorganized, with the Defence Research Group (DRG) and the Advisory Group for Aerospace Research and Development (AGARD) being merged to form a new organization, the NATO Research and Technology Organisation (RTO). Following the reorganization, RSG-30 was retained as a research Task Group (RTG) under a newly formed Information Systems Technology (IST) Panel with an interim designation of IST-005, which later was changed to IST-013/RTG-002. The title of this RTG was "Visualisation of Massive Military Datasets".

The RTG confirmed the RSG-30 interpretation that visualisation was a human activity, supported by technology, by which humans make sense of complex data. With this understanding, during the period of RSG-30 and IST-005, the Group developed the "IST-005 Visualisation Reference Model", which model has underpinned most of the Group's subsequent work under its later name "the VisTG Reference Model".

The Group sponsored and developed a workshop (IST-020/WS-002) on "Visualisation of Massive Military Datasets" which was held in Quebec, Canada in June 2000. The workshop concluded that simple displays can be useful in dealing with complex data, modular and componentware structures ease system development, user involvement needs to continue throughout the development process, support for the user's innovation, initiative, intuition, and creativity (I3C) is important especially in the face of anomalous conditions, visualisation of relationships is an important unsolved problem and evaluation is an important area of research. The importance of the interactive aspect rather than simply the presentation was brought out clearly throughout the workshop.

In December 2000, the RTG published its "Human Factors, Applications, and Technologies" report – the HAT report. This report, which included the "IST-005 Visualisation Reference Model", evaluated available technologies, applications for which visualisation technology might be useful, the probable value and difficulty of applying the technology to each application, and research requirements that promised to have the best payoff. The report was written to enable a potential user to see how existing or near-future technologies might apply to a real problem, or possibly to see that no existing technology provides a cost-effective solution. Likewise, it would enable a researcher to evaluate which research issues are key, having potentially high payoff in a number of areas, or which permit the direct possibility of implementing specific applications. The document also allows researchers and potential users in all the NATO countries to evaluate what and where work is being done, thus facilitating the development of synergistic efforts.

In 2001, IST-013/RTG-007 was superseded by a new Research and Technology Group – IST-021/RTG-007 "Multimedia Visualisation of Massive Military Datasets". The objective of this new Group was to evaluate and update visualisation systems principles, developed by its predecessor Groups, deliver a workshop in 2002 on "Multimedia Information Visualisation", the results of which would help to guide the subsequent work of the group, and consider whether it would be feasible to develop and propose a lecture series on Multimedia Information Visualisation for delivery in 2005. The Group subsequently determined that it would not be in a position to deliver an appropriate lecture series within the recommended timeframe. As a result of the successful 2002 workshop, held in Halden, Norway, and the thrusts that emerged from that workshop, the Group decided to plan a 2004 workshop, this time in Toronto, Canada on "Visualisation and the Common Operational Picture".

In the wake of the events on September 11, 2001, IST-021's parent body, the NATO Research and Technology Organisation, tasked several of its technical groups to address problems of security and defence against terror attacks. Accordingly, IST-021 requested that the Vis-NX consider the subject of information visualisation needs for intelligence and counter-terror during its 2003 workshop. This became the central theme of that workshop.

1.2 VISUALISATION NETWORK OF EXPERTS (VIS-NX)

Vis-NX was formed by the former NATO DRG RSG-30 of Panel 8 as its informal technical advisory group. It has continued to thrive under the patronage of the succeeding Research Technical Groups, currently NATO's RTO Research Technical Group IST-021, Visualisation in Massive Military Datasets. The Vis-NX supports the RTG in its mandate to study and develop methods to aid understanding and use of the contents of massive datasets. Over the years, Vis-NX has held six workshops in conjunction with meetings of the "parent" RTG, with the aim of focusing much needed research in this area.

At its inception, the Vis-NX realized a new concept in NATO research discussions and activities. In operation it offers an unofficial forum for researchers to exchange information, data and expertise. It carries some of the advantages that the NATO umbrella can offer, while avoiding some of the problems with more formal arrangements, including some Governments' occasional reluctance over the last decade to join in official arrangements.

The idea for the Network of Experts came out of the November 1994 DRG Panel 11 workshop "Visualising non-Visual Information", held in November 1994 at the Belgian Military Academy in Brussels. Some twenty five experts from several NATO countries met to discuss the state of the art in the emerging technologies supporting scientific, data, text and information visualisation, and make recommendations to Panel 11 on Panel 11's role in supporting visualisation R&D. The workshop recommended strongly that Panel 11 commission a Research Study Group (RSG) to maintain awareness and to coordinate participating Nations' work in these fields, insofar as it would relate to military and defence needs. The suggested group would operate through meetings, workshops, and by using the Internet. The workshop also recommended that the proposed RSG consider fostering an external independent network of visualisation experts that would be cognizant of, and responsive to, the needs of the RSG.

NATO DRG Panel 11 accepted the recommendation for a RSG on visualisation and endorsed the concept of a supporting unofficial network of experts. Even though the RSG was not officially formed until 1996, and then as RSG-30 under Panel 8, the Network of Experts, commonly known as "Vis-NX", had already created itself unofficially following the 1994 workshop and had established a method of working by e-mail through a mail list server run by Canada. At that time, Vis-NX already included several of the experts who have since given the forum its continuity and who have maintained the necessary effort and impetus.

Once the Vis-NX was constituted under the patronage of RSG-30, its members decided that, in addition to the normal e-mail interaction, they would endeavour to meet annually and would at the same time hold a workshop on a particular subject of interest to the RSG. The timing and location of this annual meeting and workshop would be coordinated with the RSG and would be held in conjunction with one of the semi-annual business meetings of the RSG. As much as possible, the Vis-NX would alternate its meetings between Europe and North America.

The 1994 workshop in Brussels showed the best then-available examples of data visualisation, noted enthusiastically the emerging World Wide Web as a means of communicating displays, and rapidly converged on the problem of making the display's encapsulated information and knowledge transparently available, when needed, by whom, and at the requisite granularity. There was a consensus that this was best approached with users' intimate interaction with the information corpora, with rapid and transparent turnaround of interactive feedforward and feedback. This has become a continuing preoccupation of subsequent meetings, workshops and discussions.

The first "official" Vis-NX meeting and workshop, took place June 1996 in Ottawa, Canada in conjunction with a semi-annual meeting of RSG-30. Experts, mainly from North America, met with a prime topic of massive military dataset visualisation. Among the more catholic needs and interests of information discovery, discussants began narrowing attention beyond data visualisation to information visualisation. This would become an ongoing concern for the RSG over the next few years, as it developed its own interests and definitions of task-oriented visualisation in military and civil contexts.

The May, 1997 Vis-NX meeting and workshop was held in Malvern, UK, on the DERA campus – http://visn-x.net/nx2/agenda.htm. Civilian and defence applications were shown at various points in their development [e.g. sonar and electronic warfare analyses, economic analyses]. Discussion was vigorous, centering on ways in which the Vis-NX forum could be made more useful to its participants as stakeholders. During 1997, a Vis-NX website was constructed, first at DERA, but subsequently transferred to a non-government site in Canada – http://visn-x.net.

The June 1998 meeting and workshop took place in Toronto, Canada at DCIEM. Task management and other applications were presented in a format of more extensive discussion and debate than took place at earlier meetings. Several of the presented papers were particularly relevant to considering metaphors appropriate for working visually with variable information and forms of information. This workshop included an extensive tour of the DCIEM facility.

The Vis-NX met next in Malvern, June 1999 – http://visn-x.net/nx4/agenda.htm. Again, applications were shown that related to specific military tasks. Interest developed around search engines and means to visualise text corpora. This had been an important issue for the RSG since the 1996 meeting, at least, but technology development had been slow in the intervening three years.

The Vis-NX did not meet in 2000 as the patron Research Technical Group (RTG), now designated IST-013/RTG-002 of the new Research and Technology Organisation (RTO), was itself running a formal NATO Workshop "Visualisation of Massive Military Datasets" in Quebec City that would include a number of the Vis-NX members.

In 2001, the Vis-NX met at the Virtual Reality Laboratory in Aalborg, Denmark – http://visnx.net/nx5/index.html. Discussion was focused on use of the IST-005 reference model (now called the VisTG Reference Model, which is discussed in Chapter 2) for system evaluation. The discussions were lively and spirited, and led the TG to consider system evaluation for visualisation technologies as a part of its mandate. Again at this meeting, there was a memorable facility tour of the Virtual Reality Lab.

The Vis-NX did not meet in 2002, but supported the RTG workshop "Massive Military Data Fusion and Visualisation: Users Talk to Developers" in Halden Norway. The Vis-NX adopted the format of this highly successful workshop for use in its next meeting. It chose as its workshop topic "Visualisation Needs for Counter-Terror", as requested by the RTG. The meeting and workshop were held March, 2003 at Pennsylvania State University. This was the most successful Vis-NX meeting to date, due in no small measure to the enthusiasm and momentum generated in Halden.

The Vis-NX has a select membership. It includes the members of the current patron RTG plus invited experts. Its original members were identified at the Brussels workshop of 1994 and were subsequently invited to form the Vis-NX by RSG-30 when it was created in 1996. The Vis-NX expands by inviting other experts identified by its members: Any Vis-NX member may recommend an expert for membership provided that that expert comes from a NATO or PfP country. In mid-2003, Vis-NX comprised approximately 100 members from 14 countries. The Vis-NX expects to continue to operate through e-mail and to meet annually in conjunction with an RTG meeting to discuss topics of interest to the RTG; however it does not expect to hold its own separate meeting in a year in which the RTG is sponsoring an official NATO Workshop because its members are themselves likely to play a key role in such a workshop.

1.3 NATO AND VIS-NX WORKSHOPS

Two NATO Workshops were sponsored by IST-021/RTG-007. For each, the programme committee consisted of members of IST-021/RTG-007, and their work was considered to be integral to that of the RTG. The first, was in Halden, Norway in 2002 (RTO-MP-105), the second, in Toronto, Canada in 2004. In addition, Vis-NX conducted workshops in 2001 and 2003, when there was no NATO workshop sponsored by the RTG.

Starting with the Halden Workshop, in addition to general plenary presentations and discussions of research and application, each workshop (whether NATO or Vis-NX) used several small "syndicates" or "working groups" to examine specified topics. The selected topics had in many cases been introduced by syndicates at the predecessor workshop, and were considered by the RTG to be important to develop further. For example, the Vis-NX Workshop in 2003 was asked to further develop counter-terrorism ideas reported by one of the syndicates at the Halden NATO workshop in 2002.

So far as feasible, each syndicate included representatives of the scientific research community, the developer community, and the military user community. The intention was to ensure that the ideas developed in the intensive work of the syndicate would be militarily relevant, scientifically defensible, and developmentally feasible. In most cases, this attempt was reasonably successful, though the military user community was, as might be expected, less well represented at the Vis-NX workshops than at the NATO workshops.

Communication between the military user community and the scientific researchers had been a problem at the NATO workshop in 2000 that had been sponsored by IST-013/RTG-002, but seemed to become less of a problem at the later workshops in 2002 and 2004. Many of the same researchers and several of the more senior military attended more than one workshop, which made for easier communication at the later ones, as the researchers became better aware of military needs, and the military of scientific possibilities.

Several of the same issues emerged at each workshop. One of the critical ones was the problem of displaying relationships. Almost all action requires the actor to assess relationships. Even in such a trivial everyday action as to transport an object, the relationship between the object's size and the capacity of containers, the relationship between the terrain and the available transport mechanisms, the relationship between the cost of replacing the object and the risk of loss, are only a few of the relationships that must be implicitly or explicitly considered. Furthermore, there are second-order relationships among even these simple ones. Techniques for displaying relationships in a manner that aids visualisation are not well developed.

The concept of relationships inevitably implies the concept of networks. If A, B, and C each has some kind of relationship with the others, the set of relationships forms a network in which A, B, and C are "nodes" and the relationships among them are "links". Some networks have well-defined link structure, in the sense that if A is linked to B and not to C, an action by A is propagated to B and not to C. Other networks have a less well-defined structure. For example, A may act on an environment that is available for inspection by B and C, but which will not necessarily be so inspected until much later, if ever. Yet again, some networks may be linked by momentary broadcast; A may broadcast a message, but B and C will receive it only if their receivers are turned on at that specific moment. In neither of these latter cases is the network structure defined beyond a probabilistic statement about the existence of the link.

Networks have emergent properties beyond those implied by the nature of the nodes and of the relationships among the nodes. For example, the nodes in a network may be connected randomly, they may be connected as a set of branches radiating from a single hub, they might be connected as a hierarchy of networks connected locally in random way, but with the local nets connected through higher-order networks of hubs, or they might have other statistically describable patterns of linkage. These patterns have strong effects on the ways networks behave, and affect the requirements for their display. For example, a road network includes roads with widely different traffic-carrying capacities, from multi-lane expressways to rutted cart tracks. A display of the road network for one purpose such as showing the fastest routes between major cities might include only the expressways, whereas a display for hikers might show the cart tracks and indicate the expressways only as obstructions.

When the different properties of the nodes and links in a network are added to the structural differences among networks, the display requirements become very challenging. At the 2004 Toronto Workshop, one of the working groups attempted to draft a set of abstractions of node and link properties that could be used in developing display requirements for different purposes. They then suggested which of these properties would probably be important for user purposes in different applications. As examples, they chose the application areas of counter-terrorism, information assurance, and logistical analysis.

Networks not only have static properties defined by the analysis of their node-link structure and the natures of the nodes and links, they also have dynamic properties. Events in one part of a network may propagate along the links to other parts of the network. Their influence may dissipate over time, may grow and then dissipate, may cause oscillations, or may develop chaotically. Visualisation requirements may include the provision of mechanisms for users to assess the probable future evolution of the effects of different action choices. One obvious area in which this would be important is in the delicate socio-political networks involved in peacekeeping operations.

The 2004 Toronto workshop (IST-043) had as its theme "The Common Operational Picture". It became obvious from the work of several of the working groups that the concept of a common operational picture left much to be desired. People collaborating in any venture, whether in a battlefield that contains a definable enemy, or in the design of a complex device, need to know what the other actors know and intend, but they do not need to see the same "picture". Each collaborator has a different purpose. To fulfil that purpose, she needs to know the overall objective and where the purpose fits into that objective, and needs to know where the purposes of others also fit, especially if their actions influence the environment in which she would act. The terms “Common Operational Environment” and “Common Knowledge of Intent” might be more appropriate than “Common Operational Picture”.

At every workshop, the importance of real-time interaction between the user and the display was emphasised. The ability of the user to control not only the display content, but also the manner of display (as, for example, the viewpoint in a simulated scene), greatly affects the effectiveness of the user’s visualisation. At a very early Vis-NX workshop, for example, the claim was made that in a display of the message-passing and inheritance structure of a moderately large software system, a passively rotated 3D display allowed the user to visualise about twice the amount that could be visualised from a 2D display, whereas an interactive 3D display in which the user cold manipulate the viewpoint increased the advantage to a factor of between five and seven.

As the several workshops pointed out, not only the presentation on the display surface, but also the user’s interaction with the display is likely to be critical to the success of a visualisation system. For this reason, IST-021/RTG-007, following the lead of its predecessor groups, used a reference model based around the concept of a layered structure of interaction. This “VisTG Reference Model” presented a framework within which different display technologies could be deployed, but that also suggested to designers and evaluators the possibilities for displaying not only the data to be visualised, but also the control the user should be given over the displays in support of different kinds of purpose.

1.4 THE VISTG REFERENCE MODEL

1.4.1 Introduction and Background

The VisTG Reference Model was developed initially by the predecessor of IST-021/RTG-007, initially under the name of the “IST-005 Reference Model”. Under that name, it was described in the report RTO-TR-030, “Visualisation of Massive Military Datasets: Human Factors, Applications and Technologies”, colloquially known as “the HAT Report.” The Model has provided one of the intellectual bases for consideration of visualisation by IST-021/RTG-007, whose mandate was primarily concerned with methods for the evaluation of visualisation systems. Much of the work of the group was based on the idea that the VisTG Reference Model could be used in evaluating the effectiveness of visualisation systems, at least to the extent of identifying possible alterations that would be likely to improve their effectiveness.

The predecessors of IST-021/RTG-007 recognized that visualisation systems are quite diverse in their design and mission, so that to evaluate different visualisation techniques and to compare systems is difficult. The VisTG Reference Model is intended to provide a framework within which both evaluation and design can be performed, no matter what the mission or complexity of the system in question, and no matter what software and hardware is used. During the life of IST-021/RTG-007, the use of the Reference Model for design was addressed only incidentally, when evaluation suggested modifications to some design.

The VisTG Reference Model depends on answering the question: “What is the purpose of visualisation?” That question is the common thread that links all questions of the form “Why is visualisation needed in this instance?” This one question, in any specific instance, induces a family of further questions related to the situation at hand. In the general case, the initial question induces a family of further questions that together determine the Model’s structure.

Visualisation is taken to mean something that happens in the head of a human that forms one route toward understanding some structure or situation. It is related to the concepts of “intuition”, “seeing at a glance”, or "getting the picture". Another component of understanding is logical analysis, for which a colloquial concept might be "working it out".

When one is "working it out", one is concerned at any moment with the interrelationships of very few well-defined entities that influence each other and that jointly affect the output of some "working out" algorithm. In contrast, when one is "seeing at a glance", one sees the coherent effect of many elements that are not necessarily perceived individually, and one may not even be able to tell the reason why a particular pattern or signification has come into one's head. Nevertheless, it is often possible to check the implications of visualised situations through logical analysis. The reverse also holds; the results of logical analysis may often be visualised. Logical analysis without visualisation, based on the relationships among isolated data, can lead to conclusions that, though logical, are wildly wrong, as when a (logical, analytic) bank computer shows that one has recently made a deposit of $10,000,135.36 instead of $135.36. In such cases, visualisation can sometimes very quickly show that the analysis must be faulty, or be based on faulty premises.

The VisTG Reference Model explicitly ignores the fact that visualisation complements and is supported by analysis in the development of human understanding. Displays that best support precise logical analysis may be different from those that best support visualisation. Displays for logical analysis tend to allow the human to focus on individual items and relationships, whereas displays that support visualisation tend to emphasise context and pattern. Both are usually necessary, as logical analysis by itself often results in locally but not globally optimum behaviour, whereas visualisation by itself often provides imprecise, but globally reasonable interpretations to guide adequate, but not necessarily optimum, behaviour. One purpose of visualisation, then, is to provide a large context within which the sense and import of the results of logical analysis can be made effective.

A question that arises immediately is the purpose of providing context for the meaning of logical analysis in a particular case. The purpose, presumably, has to do with the intent of the person concerned. A field commander may be planning an attack or a defence, or simply trying to assess the intentions of an enemy. A software designer may be considering the interplay of messages within a complex system. A Homeland Security analyst may be looking for dangerous points of possible failure in the national infrastructure. An intelligence officer may be looking for the patterns of communication that suggest the formation of an active terrorist cell. In each case, the final intention is to do something about the world being visualised and analyzed. The field commander visualises the playing out of his and the enemy's intentions; the software designer visualises opportunities and problems in the synergism and interference of the messages being passed; the Homeland Security analyst visualises the effect of a destroyed power distribution station on the water supply to a city, and the consequent civic unrest. The intelligence officer visualises the new cell and its probable future development and activity. In each case, visualisation allows the person to contemplate beneficial alterations to the situation visualised.

People visualised situations long before computers existed. Alexander, Caesar, or Napoleon would have visualised the next day's battle before issuing the battle orders, and would have visualised the ongoing battle as messengers brought fragmentary reports to augment whatever he might have been able to see from his hilltop vantage. In the modern environment, many of those reports would be collected in some computer system (and there would be far more, and more complex, reports to aid and to confuse the commander). The commander would then gather the collected information from some kind of computer display, be it visual, auditory, or haptic, pictorial, diagrammatic, or in words.

The computer display ordinarily would provide only one of a number of inputs to the commander's visualisation. Evaluations, and even off-the-cuff comments and questions from specialized staff also contribute, but those officers also are very likely to have gained their information, and to have created the visualisations in their heads, from data saved in some computer's dataspace. Comparable considerations apply to any other person who uses a computer as an aid to visualising a situation of interest. For much of the following, we will simply talk about "the user" rather than "the commander, software analyst, security analyst, ...".

1.4.2 Definition of a "Visualisation System"

A definition of a "Visualisation system" used by IST-021/RTG-007 is: *A system for presenting, probably interactively, some part of a dataspace, in such a way that a user with some purpose in mind can visualise the import of the data for that purpose*.

This definition has several independently important components:

- **Presenting**: The system displays the data appropriately for the user's senses, which does not necessarily mean that the display is pictorial.
- **Interactively**: The user can influence what the system presents and how it presents it, in real time.
- **Some part of**: Not all the data can be displayed simultaneously.
- **Dataspace**: At any one moment, a delimited set of data is available for display, though more may be being acquired on an ongoing basis.
- **Purpose**: The user is trying to perform a task for which the system may be of some assistance.
- **The import for the purpose**: To visualise the data is not the user's objective. It is only a means to an end.

The most important of these elements is probably the last. Visualisation and the use of a visualisation system is only a means to an end, not the end in itself. The very best presentation of the data will be useless if the data, as interpreted, are not relevant to the user's purpose.

Purposes, however, change on various time scales. At the long end of the time scale the user may have an objective that takes years to achieve. At the short end, the user may purpose just to scroll to the next segment of a map, or even to hit "y" on a keyboard. Purposes on all time scales are held, and acted upon, simultaneously. Those involving short time scales are often the means by which the longer-term purposes are accomplished. The VisTG Reference Model provides a way to integrate them all into a common framework.

1.4.3 Development of the Basic VisTG Reference Model

The VisTG Reference Model, deals with visualisation based largely on data held in the dataspace of some computer, for the purpose of allowing a user to influence events in the outer world. To acquire useful data from the dataspace, the user must be able to query the dataspace so as to have it display data relevant to the user's purpose, and there must be a way for the dataspace to provide the desired data. Typically, the data first requested will not be ideally suited to the user's requirements, so the user must be able to influence further the data to be provided. A feedback loop must exist between the user and the dataspace. This feedback loop is the seed concept for the VisTG Reference Model, as shown in Figure 1-1.

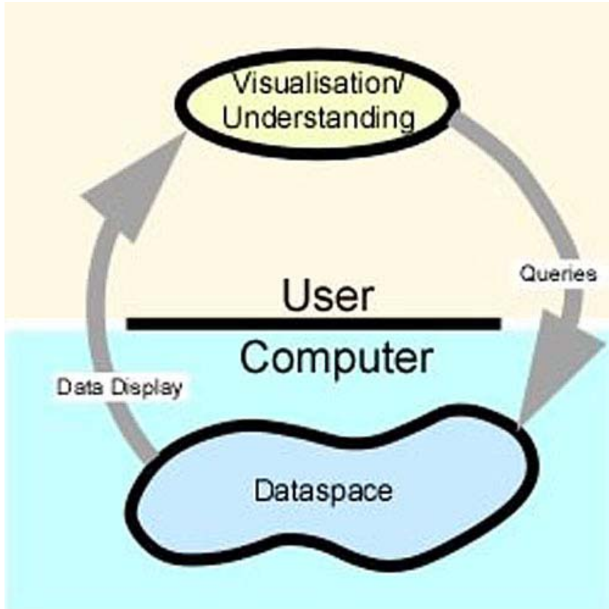

Figure 1-1: The First Stage in Developing the VisTG Reference Model, Showing the Feedback Loop between the Users' Visualisation and the Dataspace Inside the Computer.

Figure 1-1 is derived from two statements:

- *The user must be able to inform the dataspace what kinds of data are required*, and
- *The dataspace must be able to deliver those data to the user in a way that facilitates visualisation and understanding*.

Each of those two statements immediately invites the question: "How is that done?"

Obviously the user has no telepathic connection with the dataspace. Something must translate the user's queries and thereby invoke procedures that select and process the desired data. We give those intermediary systems the generic name "Engines." Engines are also required to convert the raw data extracted from the dataspace into a form of which the user can take advantage.

On the human side, as noted above, visualisation is not the same as understanding. It is one of the two major routes toward understanding. What the Engines provide is what allows the user to visualise, and what the user wants to visualise determines how the queries are presented to the Engines. Figure 1-1 therefore needs to be modified on the human side as well as on the computer side.

These two considerations imply that the reference model needs an inner structure. We still keep the connecting loop between the human "understanding" and the computer "dataspace", but we add an inner loop that expresses how the outer loop is actually implemented. Figure 1-2 shows that what the user wants to achieve is "Understanding" the dataspace, and "Visualising" is a means to that end. The user's understanding depends on interaction with the dataspace, but is (in part) implemented through visualisation, which requires interaction with the Engines that interpret the user's requirements, communicate with the dataspace, and manipulate the data for presentation.

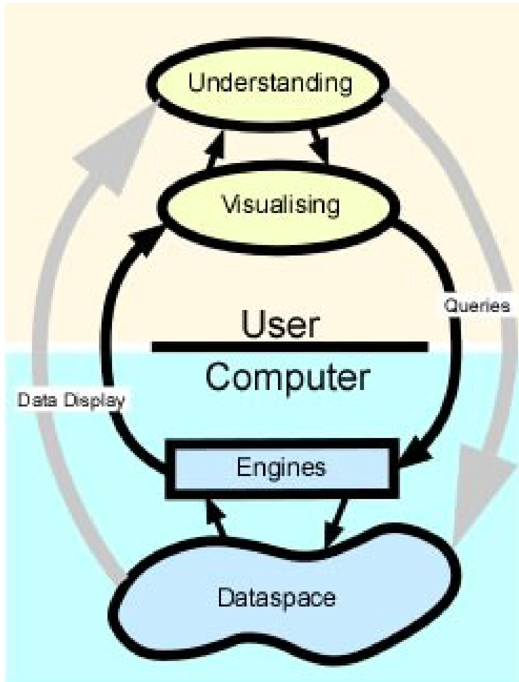

Figure 1-2: First Elaboration of the VisTG Reference Model, Showing that the Communication Between the Database and the User's Understanding is not Direct, but is Implemented Through Visualisation and the use of Engines that Interact with the Dataspace.

Engines are of many different kinds, but there are two main classes:

1) Those that interact with data from the dataspace and manipulate them in some way, perhaps adding the results to the data space, perhaps displaying the results to the user, and
2) Those that do not interact directly with the data, but work with the user in determining how the data should be selected or manipulated.

Because the user needs to control the actions of the Engines, each individual Engine must be involved in its own feedback loop with the user. Each single arrow in Figure 1-2 represents many parallel threads. The user must be able to understand what an Engine can do and is doing, which may involve the user analysing or visualising the Engine's behaviour. The user must be able to instruct each Engine, and the Engine must be able to display to the user the necessary information that permits the user to determine how those instructions actually affect the actions of the Engine. Conventionally, to keep the picture simple, these individual user-to-Engine loops are omitted from diagrams of the VisTG Reference Model, but they must be considered when using the Model for system design or evaluation.

Telepathic communication with the Engines is no more possible than is telepathic communication with the dataspace. Communication must be through physical devices and their associated software. This implies the addition of a third, innermost, set of communication pathways, as shown passing through the input-output devices in Figure 1-3. Conceptually, the feedback loops conform to the grey arrows, but the actual implementation follows the black arrows in Figure 1-3. Understanding has links to the dataspace; visualisation has links to the Engines, and all the links are implemented through the physical devices and their support software.

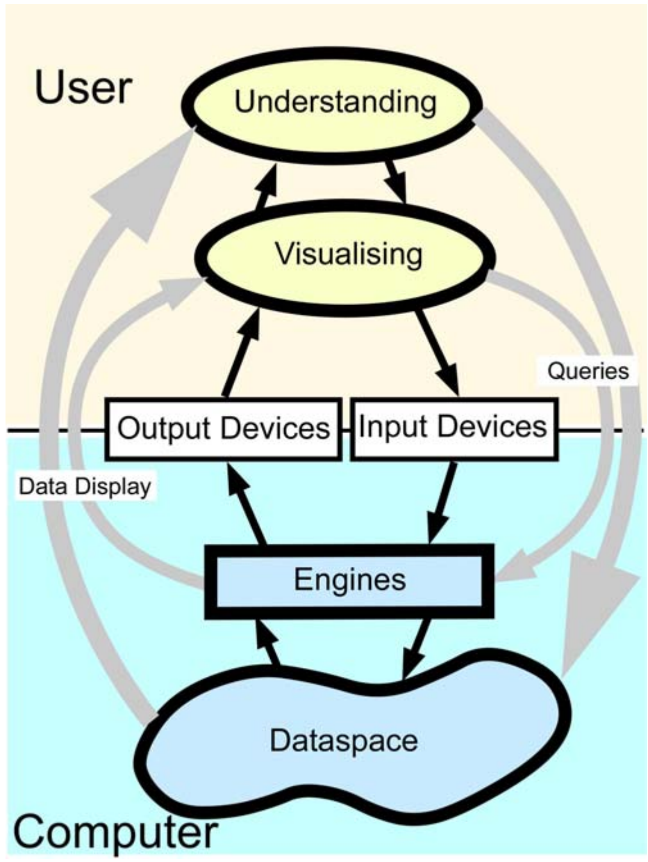

Figure 1-3: The Final Elaboration of the VisTG Reference Model, Showing the Physical Devices that Enable Communication Between the User and the Processes and Data in the Computer.

Figure 1-3 is the canonical representation of the VisTG Reference Model. The user's primary objective is to understand something in order to be able to act appropriately. Data helpful to that understanding may exist in the computer. We assume that it is psychologically correct that in most humans, visualisation is an important route to understanding, and that current understanding of a situation influences how one visualises the import of new data. In other words, one's perception depends in part on what one knows or expects. We acknowledge that there is a complementary route to understanding through logical analysis, and that this other route is equally important, but we do not include it in the canonical reference model for visualisation.

On the other side, within the computer, we consider that there is a dataspace, which consists of all the data in the computer that might possibly be of interest to the user, and much else, besides. The data in the dataspace are, however, of no value to the user unless they are interpreted and massaged in ways appropriate to the user's purposes. Interpretation and massaging of the data is the job of the Engines. Just as there is a feedback loop within the human between understanding and visualising, so there is a related feedback loop in the computer between the Engines and the dataspace.

Finally, the human and the computer interact through the input and output devices that both can observe or affect. The Engines that massage the data can be instructed as to what the user wants, and can display the results to the user, through those devices and their associated software. Likewise, the user translates what is displayed into visualisations using the same facilities normally used for interpreting the natural world, and informs the Engines using the same facilities used for communication with other people, though, of course, with different emphases necessitated by the different interpretive capacities of computers and people.

1.4.4 Refinement of the VisTG Reference Model

The canonical diagram of the VisTG Reference Model shows the categorical structure of the feedback loops at the three major levels of abstraction. These, however, are not the "atoms" from which the Reference Model is constructed. They are the conceptual shape of the structure formed from the interrelationships of the actual functioning feedback loops. Each of the canonical feedback loops is formed from possibly very many individual feedback loops between the user and the computer, many of which are supported by yet other loops. Each of these "atomic" loops is determined by a single well-defined purpose or objective. Its actions, if successful, allow the user to achieve that objective. The actual structure and interrelationships of these loops depends on the system being analysed or designed, and on the user's purpose, training, and situation awareness. The way the "atomic" loops function can be described for the general case, but the details of their interrelationships is different in each specific case.

An "atomic" loop superficially resembles the well-known OODA (Observe, Orient, Decide, Act) loop. It differs, however, from the OODA loop in a very important way. Whereas the OODA loop is usually conceived as starting with "Observe", the atomic loop of the Reference Model starts with the user's objective, which is outside the loop. Observation is done in order to see how the current situation fails to meet that objective. The user perceives a situation in the environment (in this case, perhaps the dataspace, perhaps the state of an Engine or the brightness of a screen) compares it with the desired objective, and acts on the environment, which changes the state. The critical point is that the existing condition in the environment is compared with an objective for that condition, and the actions are expected to reduce the difference between the actual and the objective condition. The objective, in this simple loop, is the "purpose" of the action. The basic loop is shown in Figure 1-4. Figure 1-4 also shows some of the potentially important problems that must be addressed in the evaluation of any such loop.

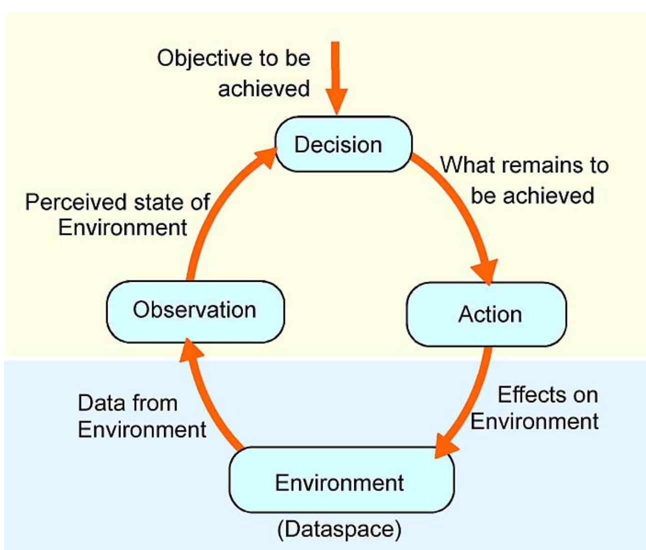

Figure 1-4: The Basic Control Loop on which the VisTG Reference Model is Built, Showing Some Things that Should be Considered in Design or Evaluation.

Many such atomic loops work together, just as physical atoms work together to make chemical compounds. As an example, imagine a briefing session over a map display. A certain segment of the region of interest is already being displayed, but the briefing officer wants to show the disposition of certain forces not currently on the map. To do this, the briefing officer must be able to select the forces in question from among all the data held in the computer, to select which attributes of those forces should be represented on the display, and to ensure that the region of the map being shown is appropriate to the disposition of these particular forces and their mission (if friendly) or potential danger (if hostile).

At one level, the objective is that the map should show the correct region, and on it the force dispositions should be shown with the right level of detail to make the point the briefing officer wants to get across. At another level, there are several objectives, such as to inform the computer as to the appropriate selection criteria for the forces that are to be displayed, to ensure that the correct map region is selected for display, to ensure that the right geographic features are selected for map display, and to ensure that the force display contains only and all the attributes required for the briefing purpose. Each of these objectives, and others not mentioned, imply the existence of a loop of the kind shown in Figure 1-4. However, for these loops, the objectives are determined by the main objective of presenting the desired information in the briefing.

The subordinate loops are not simply dependent on the purpose of the main loop. They implement the action component of the main loop, as shown in Figure 1-5. Furthermore, the observations that form part of their individual loops all may contribute to the observation that is compared with the objective of the main loop. Each of the user's purposes implies the existence of a loop.

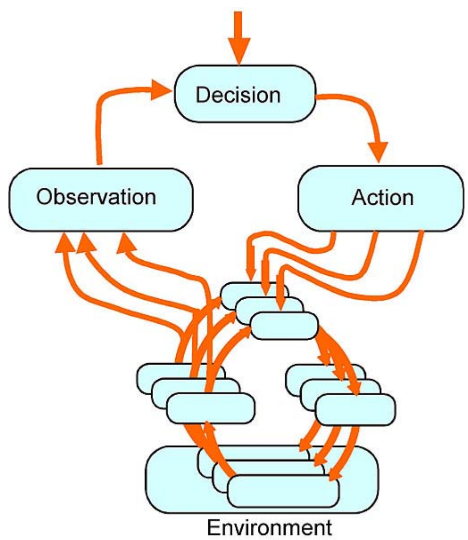

Figure 1-5: The Actions of a Loop are often Performed by Subordinate Loops whose Objectives are Supplied by the Main Loop. This kind of subordination can be repeated.

Sometimes we may identify part of a loop or of a parallel set of loops with a single element of the Reference Model, such as an Engine. Engines exist to serve purposes of the user. Typically, a task-level purpose is likely to take advantage of four kinds of Engine:

1) **Navigation Engine**: Navigate to the relevant part of the dataspace.
2) **Data Selection Engine**: Select the relevant data.

3) **Algorithm Selection Engine**: Choose algorithms to apply to the selected data.

4) **Algorithmic Engines**: Execute the appropriate algorithms on the selected data.

Among the “algorithmic Engines” may be Engines that prepare the manipulated data for display, Engines that alter the data in the dataspace, Engines that affect the operation of the computer-user interface, and an indefinite variety of others. Which of these may be appropriate to the situation depends on the user’s main purpose and on the nature of the data on which the algorithmic Engine may be used. This multiplicity of possibilities is the reason an Algorithmic Selection Engine is often required.

A basic loop never operates in a vacuum. There are all sorts of reasons why one or other of the links in the loop might be problematic as illustrated in Figure 1-6. There could be intrinsic reasons why the user might not be able to perceive data that would be available in the database. For example, the user’s training might not allow the data to be interpreted in a mission-sensitive way. This would be a “possible impediment to observation” in Figure 1-6. An “external influence” might be a requirement for the user to pay attention to some other data while working on this specific loop. Similarly, there might be internal impediments or external influences affecting the user’s ability to influence the environment (meaning, for our purposes, the Engine with which the user is interacting, or the dataspace). When evaluating a system, these possibilities must be considered.

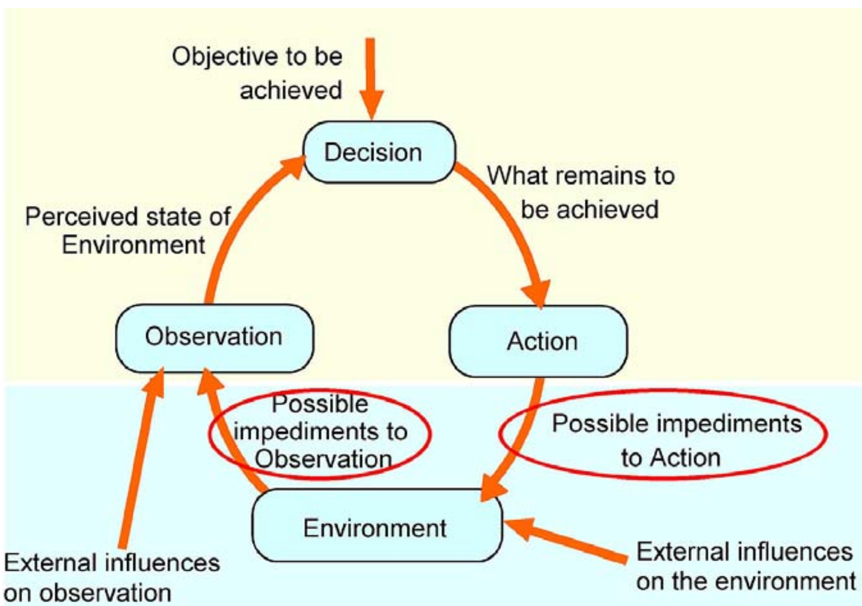

Figure 1-6: Some Influences that may Affect the Performance of the Loop. They must be considered when evaluating the system that contains the loop.

1.4.5 Using the VisTG Reference Model

The VisTG Reference Model is intended to provide a framework within which both evaluation and design can be performed, no matter what the mission of the system in question, and no matter what software and hardware is used. For the purposes of IST-021/RTG-007, the Model was used to provide a test of a possible method of evaluating visualisation systems.

One has to ask what is meant by "to evaluate." Using the thinking that led to the Reference Model, "to evaluate" means to assess how well the system allows a user to accomplish a purpose, and to propose design changes or changes in user training that might improve the situation. A hypothetical example of how The VisTG Reference Model may suggest design improvements was provided in a presentation at the IST-036/RWS-005 workshop (Microsoft PowerPoint at ftp://ftp.rta.nato.int//PubFullText/RTO/MP/RTO-MP-105/MP-105-S2-03.pps). Additionally, "to evaluate" may well mean that this initial assessment is compared to how well potential variants or other real or hypothetical systems might serve the user.

It follows that any evaluation using the Reference Model must start by determining a purpose. The purpose suggests the existence of an atomic loop, and the evaluation considers whether the implied loop exists, and if it exists, how well its elements serve the purpose. Within the Model, an atomic loop might be between the human's "understanding" and the dataspace, or between the human's "Visualisation" and an Engine. Any realistic system is likely to have many loops of each kind.

Once a single loop has been defined, questions arise about its operation and about its relationships with other parts of the system.

"Why might this purpose exist?" is a question about higher-level purposes, and therefore about atomic loops for which this one serves as part of the means of action (as shown in Figure 1-5).

"How are the observation and action processes implemented?" is a question that may well imply the existence of supporting purposes, and hence subordinate atomic loops.

"What effects might other notionally independent loops have on this one?" is a question about possible interference or support provided by other parallel loops. For example, to fill a screen with text makes it hard to read a map on the same screen, and vice-versa. However, this question can be answered only after other loops have been defined (Figure 1-7).

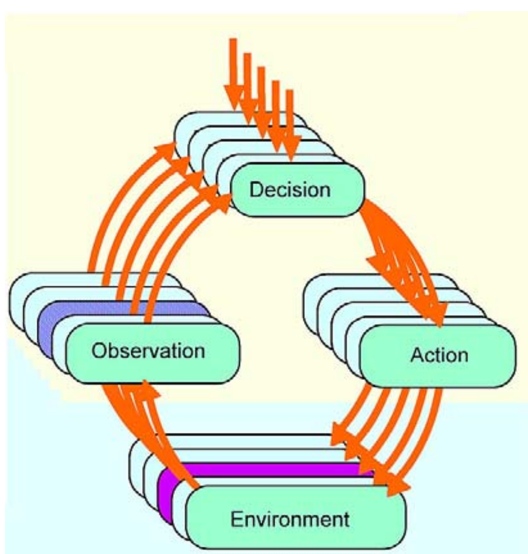

Figure 1-7: Many Loops are Probably Potentially Operating in Parallel, but the User Cannot be Attending to all of them. Is there a way the system can alert the user to the possibility that something important may be perceptible in a currently unattended loop?

"What effects might this loop have on others?" is a parallel question that should be answered if other notionally independent loops have already been defined. Implicit in these last two questions is the fact that the number of bilateral relationships varies as the square of the number of entities, which further implies that the effects of C on B may alter a preliminary analysis of the effects of B on A. This explosive proliferation of potential effects may well be the most difficult feature of any evaluation that goes deeper than the user's overall satisfaction and performance, whether based on the VisTG Reference Model or not.

There is another effect that other loops might have on the one being analysed. Clearly, in most dataspaces, too much is going on for the user to take it all in at once, either because there are incoming data streams not under the control of the user, or because the dataspace has such structural complexity as to make it likely that the user may miss significant patterns unless the possibility that such a pattern exists is drawn to the user's attention. In other words, something may occur or exist that demands the user's attention in the part of the environment that is of concern in a presently unattended loop. The question then arises as to whether there is a means, within the loop being analyzed, to alert the user that something potentially important may be worth attending to elsewhere.

These issues, and others such as the importance of user experience and training, are made explicit in a series of questions that should be asked, starting with any purpose that can reasonably be asserted for the intended or actual user of a visualisation system. The answers to these questions often imply the existence of other atomic loops at higher and lower levels, about which the same questions may be asked, until the entire workings of the system have been analyzed.

1) What is the user's purpose that defines this loop?
2) How well can the user perceive what is necessary in order for the user to achieve the objective?
3) How well can the user choose the necessary actions to affect progress toward the objective?
4) What impediments or external influences may affect the user's ability to perceive what is necessary, and what are the implications for training?
5) What impediments or external influences may affect the user's ability to act effectively, including implications for training?
6) How may other parallel events or tasks alert the user that it might be valuable for the user to attend to them instead of this task?
7) Apply questions 1 to 6 to any parallel loops and recursively to lower-level or higher level loops whose existence is implied by the results of the earlier analyses.

It is difficult to go deeper into how these questions are instantiated without referring to a specific example, but each must be instantiated before it can be answered. Chapter 2 goes into these questions in more detail, and discusses their use in connection with the evaluation of a system called VITA. Here, then, is a much simplified version.

1.4.6 A Trivial Example of how the Six Questions Might be Instantiated

Let us imagine that the user's purpose is to visualise the route to a particular address, using an on-line map that allows for changes in the area displayed and the scale of the display, such as "MapQuest" (http://www.mapquest.com). Here is a sketch of the analysis of a single loop.

1. *Describe some purpose of the user.*

   The user wants to see the street plan in the neighbourhood of the target address.

2. *Examine the degree to which the user can perceive what is necessary in order for the user to achieve the objective.*

   Initially, the map system is likely to display a map of a random location at a random scale. The question to ask is whether there is a means for the user to relocate the area displayed so that it includes the target address at a scale that allows the user to ascertain the route. This question implies further questions about the capability of the Engines with which the user must communicate. Clearly, a new purpose, and therefore another loop, is implied: the new purpose is for the user to bring the map display to a useful location and scale. Assuming that is possible, then the question at the current level can be addressed: Assuming that the map is displayed so that it shows the target address and the street plan around it at an appropriate scale, is it displayed in a way that allows the user to visualise a route? There are several ways in which the answer might be negative. The map might lack street names, or might show them in letters too small for the user to read. It might not show one-way streets in a way that the user can understand. It might not show the differences between major marked routes and side streets, or which streets are likely to be hard to negotiate because of on-street parking.

   If the answer to this question is that the user can visualise a route, but might not be able to follow it in practice, because the map display failed to show advantages and difficulties of various possible routes, we are actually answering a question about a higher-level purpose: Perhaps the higher-level purpose is that the user wants to go to the address in question, or perhaps it is that she wants to redevelop the one-way pattern of streets to ease traffic flow, to organize a plan for a robbery, or to visualise how to deal with terrorists holed up at the target address. So, these answers imply the existence of a higher-level loop, one supported by using the on-line map. That higher-level purpose also refines the purpose we are now considering, in the sense that the features shown on the map may or may not be part of what the user "wants to see" in the street plan.

3. *Examine the degree to which the user can choose the necessary actions to affect progress toward the objective.*

   Progress, in this case, is bringing the map nearer to the target location and scale, and, if necessary, causing the desired features, such as street names or one-way routes, to be displayed. At this point, the question is not how the user can do this, but whether it can be done. If the Engines provide the means, it is reasonable at this point to ask whether those means are effective or difficult to use, but the "How" question invokes another set of supporting loops, some of which may already have been implied by the answer to the previous question. The analysis of those loops may be the more appropriate time to deal with why such problems might exist.

4. *Examine any impediments or external influences that affect the user's ability to perceive what is necessary, including implications for training.*

   In the example at hand, it is hard to see many potential impediments. We have to assume that the user knows how to read road maps, because if this were not the case, she would never have resorted to the map display in order to discover her route. However, we can imagine impediments imposed by poor map display, such as misplaced or missing street names. We can also imagine impediments based on how much or little the user knows about reaching the general area of the target. It might not be possible for the map display to show a route that starts in the user's area of knowledge and at the same time shows the street plan around the target. That kind of impediment is a crude example of taking into account the user's training. In other situations, however, one can easily imagine impediments such as simultaneous requirements of other loops for the same screen real estate, though this is a problem to be noted here, and examined when the time comes to analyze the I/O device level of the system. Here, the equivalent impediment might be a requirement on the user to pay attention to two incompatible things at once.

5. *Examine any impediments or external influences the affect the user's ability to act effectively, including implications for training.*

   The usual kind of impediment at this level is again likely to be a conflict with the actions required by another loop. Perhaps the user's attention might be needed by the other loop, perhaps the other loop requires incompatible manipulations of the I/O Devices (which, however, should be analyzed at the I/O level rather than here). Here again, the potential problem may be noted for later examination, when the displays and input devices are under consideration.

6. *Examine the degree to which other parallel events or tasks alert the user that it might be valuable for the user to attend to them instead of this task.*

   Whether this question is applicable depends on other parallel purposes have been identified or are anticipated in the analysis of the larger system. When one is dealing with a single purpose, as in our example of the on-line route map, there is not much opportunity for alerting. However, in the analysis of a more complex system, particularly one in which information in the dataspace changes independently of the user's actions, opportunities for alerting may well arise.

7. *Apply steps 1 to 6 to any parallel loops and recursively to lower-level or higher level loops whose existence is implied by the results of the earlier analyses.*

   The analysis at this point has implied that there must at least be loops for which the purposes are to set the location and scale of the map display. These are lower-level loops. It has also implied the possibility that there might be lower-level loops to define the features displayed on the map, and at least one higher-level loop defined by the purpose for which the map display might be wanted.

   This latter implication of a higher-level loop is quite important when one is evaluating a tool such as VITA, which is discussed in Chapter 2. A few widely different higher-level purposes for the on-line map were mentioned above, and these carry different implications for the requirements on lower-level purposes that would need to be investigated in step 7.

   This example is perhaps too trivial to illuminate the possible ramifications of the questions, but it does illustrate one of the problems that tends to arise. In steps 2, 3, 4, and 5, the evaluator is very likely to have in mind not only the answer to the question, but also the reason for the answer. The reason for the answer usually is a statement about the behaviour of some supporting loop, and when the analysis comes to that loop, the reason may be an answer to one of the six questions, with its own reason. The evaluator will have made the error of "mixing levels."

   In the example at hand, If question 2 (can the user perceive what is necessary) should be answered "No", and the answer actually given is "No, because the wrong part of the map is probably being displayed", the evaluator must be careful to separate the answer into one part that deals with the question as asked, and another part that says "There should be a way for the user to alter the region shown on the map. Check whether a loop with this purpose exists, and how well it works." If no such loop exists between the user and an appropriate Engine, the evaluator notes a recommendation for a design modification (of course, in this specific case, the whole rationale for the system would fail if no such supporting loop existed, but the implications are more subtle in most realistic cases).

The VisTG Reference Model has three levels of loop. The outer is the one that concerns the task of the user. The second level concerns the mechanisms, known in the model as the loop that connects the user's ability to visualise with Engines in the computer. Generically, any visualisation system has four kinds of Engine loops at level 2:

- Navigation Engines, which allow the user to move around the dataspace;
- Data Selection Engines, which allow the user to choose subsets of the data for manipulation (such Engines are closely related to Navigation Engines, and in some systems may be one and the same);
- Algorithm Selection Engines, which allow the user to determine how the selected data subsets are to be manipulated; and
- Algorithm Execution Engines, which do the actual manipulation of the data, including the preparation of presentations such as 2D and 3D displays, textual and tabular representations, or auditory or haptic displays.

In a detailed design, each of the six questions should be addressed for each Engine in each of the four classes of engines.

1.5 SUMMARY

IST-021/RTG-007 represents the continuing requirement for the military to be able to visualise the implications of the ever-increasing amounts of data that modern technology makes available. Predecessor Task Groups (Research Study Groups under NATO Defence Research Group) had been concerned first with the visualisation of textual data, then with the problems imposed simply by massive amounts of data, and most recently with the evaluation of systems developed to aid personnel in visualising what their data means to their tasks.

The RTG and its immediate predecessor (IST-013/RTG-002) proposed workshops in support of its mission in 2000 (Valcartier, Quebec, Canada), 2002 (Halden, Norway), and 2004 (Toronto, Canada). The proceedings of these workshops are available from RTA.

In support of the work of the Task Group, a "Network of Experts" was developed. In this context, a Network of Experts (NX) is a voluntary organization operating with the patronage and moral support of the RTG, with a membership composed of researchers proposed by pre-existing members of the NX. The NX has held six workshops, usually on alternating sides of the Atlantic, the most recent having been in the USA in 2003. NX workshops are not held in years when the RTG sponsors NATO Workshops. Since 2002, the NX Workshops and the NATO workshops have followed a format in which small working groups have addressed specific key problems, often suggested by the results of the working groups in the previous workshop.

Under the predecessor group IST-013/RTG-002 and its antecedents, a model for designing and evaluating visualisation systems was developed, called "the VisTG Reference Model." It was intended that this model be the basis for tests in which different visualisation systems would be examined both inside the nation that developed them and in other nations. Several different systems were proposed by the nations, but for individually different reasons, most were never made available for testing. Two systems were examined, but neither in great detail, and each by only one investigator. In the following chapter the results of those examinations are discussed, and the VisTG Reference Model described in more detail.

Chapter 2 – EVALUATION USING THE VISTG REFERENCE MODEL

M.M. Taylor
Martin Taylor Consulting, Canada

The VisTG Reference Model was proposed as a reference for both designers and evaluators of visualisation systems. It was intended to serve as a guide to asking questions that might lead to improvements in existing systems and that might avoid potential problems in systems being developed. IST-021 intended to use it, as well as an evaluation guide based on information-theoretic principles, to examine several different visualisation systems that were offered by the member nations. In the event, for different reasons, most of the systems proposed for evaluation were not made available to IST-021. This chapter, then, is based largely on a short attempt to use both techniques to do a quick survey of the Norwegian NORCCIS II system, and a rather longer examination of an early version of the Canadian VITA system using the VisTG Reference Model.

2.1 THE VISTG REFERENCE MODEL

The VisTG Reference Model is based on the fact that what matters in visualisation is what happens in the user's head. What happens on a screen is only part of the route toward the user's developing an understanding of the data being displayed. People do things for some purpose, including visualising the import of the data.

The VisTG Reference Model was introduced and described briefly in Chapter 1. Here, that description is fleshed out, and its use with a real visualisation system is described. Some repetition of material in Chapter 1 is inevitable, and may be skipped.

The VisTG Reference Model starts with the assumption that the user wants to visualise something that depends on the content of a dataspace within the computer. Since the user's mind has no direct access to the dataspace, the data must be first transformed into a form accessible to the user's senses. This form is typically visual, though other modalities can be helpful. Likewise, the computer, lacking direct access to the user's mind, needs interface devices through which the user can provide information about what aspects of the data are wanted on the display devices.

Between the dataspace and the display device, and between the input devices and the dataspace, there is a large gap. The VisTG Reference Model treats all the algorithms that convert keystrokes, speech waveforms, pointer movements and the like into instructions for data selection and transformation as "Engines". Similarly, Engines perform all the computation that transforms selections from dataspace into displayed images or sounds.

In the VisTG Reference Model (Figure 2-1), there are thus three major loops that represent the interactions between the user and the dataspace:

- An outer loop, representing the user's understanding interacting with the dataspace; this loop is implemented by
- a middle loop, representing the user's visualisation processes interacting with the Engines, by means of
- physical Input-Output (I/O) devices such as screens, loudspeakers, keyboards, microphones, pointers, mice, etc.

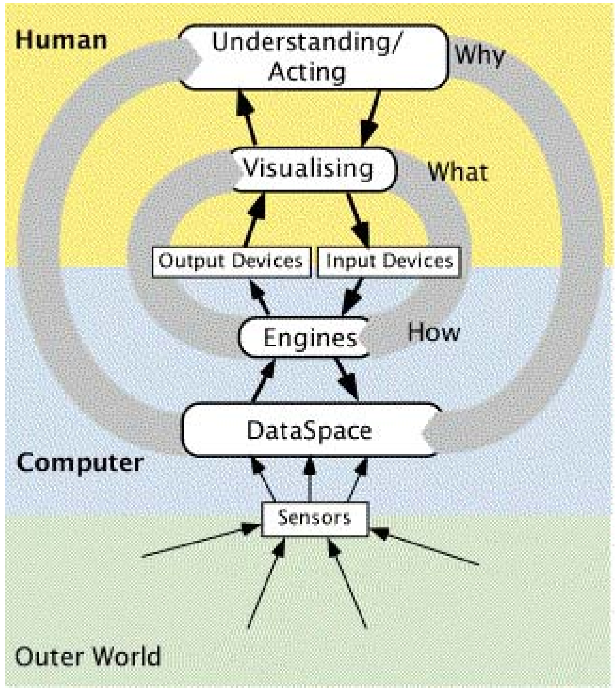

Figure 2-1: The VisTG Reference Model Set in an Environment in which the Computer Acquires Data Autonomously from an Outer World of Interest to the Human User. The conceptual information flows follow the thick grey arrow loops, but the physical information flows must follow the paths marked with small black arrows. When applied to a real system, each arrow may represent multiple parallel paths and possibly many levels of structure.

Another way of thinking about it is that the outer loop expresses *why* the user is using the system, the middle loop expresses *what* the user is doing with the system, and the innermost operations (I/O) express *how* the user is doing it.

Everything starts from the user's purpose – the "why".

In Figure 2-1, all the arrows represent types of connection for information-passing. Each should be considered as representing many paths in parallel. In a particular system some of those paths may be complex and many-layered. This is especially true of the loop "Visualising to Engines and back". The following sections describe the components of the VisTG Reference Model and their use in evaluation in more detail.

2.1.1 Understanding and the Dataspace

To visualise something is not usually the user's purpose. Visualisation is a mental tool, which works hand in hand with analysis to create understanding. Although a flow diagram like the VisTG Reference Model could be drawn to represent analysis and understanding, analysis works in the head very differently from visualisation. Effective analysis depends on the user (or the computer) performing defined operations on defined entities. Humans are not very good at this, at least not when compared to computers appropriately programmed. Humans have a hard time dealing with more than single-digit numbers of individual entities. If presented with many similar entities on a single screen, humans have a tendency to be overwhelmed, and to call it "information overload", "dataflood", "clutter", or the like.

In contrast to analysis, visualisation depends on the interrelationships among many entities, most of which are not perceived as individuals. It is rare for "clutter" or "information overload" to be a problem for visualisation, even when the self-same display could be rather too cluttered for effective analysis. Visualisation could also depend on the relative movements of a few entities over time, but what it does not depend on is algorithmic operations on a small number of individuated entities. This difference between analysis and visualisation is often critical in considering the effectiveness of display techniques, and therefore in considering the interactions between the user and the Engines in the middle loop of the VisTG Reference Model.

The user wants to understand the dataspace for some purpose, the "task". The task may involve acting on the dataspace or it may be concerned with the world outside the computer. Whichever it is, the evaluator must be concerned first and foremost with whether the system assists the user with the task, not whether the user finds it pleasing, or whether the individual elements conform to formal guidelines.

Having said that, it is very often the case that a system that pleases the user and that conforms to guidelines will also be one that helps the user with at least some tasks. Furthermore, any reasonably complex system will be used for many tasks, often several simultaneously. To assess how well a completed system serves a user for a task requires a reference standard, which could be how well the user can perform the task without the system, or whether a different or a modified system would serve the user's purpose better. Except in trivialized situations, it is virtually impossible to assess the value of a system for a task in any meaningful metric. How good is a surveillance system that misses one percent of the targets against which it is supposed to provide warning? If the alternative is that 50% would get through, it is very good indeed, but if a different and equally costly system would miss only 0.01%, then it is a very poor choice.

It is in this context that the VisTG Reference Model is intended to aid evaluation. It does not lay claim to measuring the usefulness of any system, but it does claim to help an evaluator to determine wherein complex systems might be improved, or, used prospectively, to help a designer avoid some pitfalls that might be detrimental to the value of a system under development. It does this by leading the evaluator to examine the interactions among the many different purposes of a user controlling a complex system, and the actions performed by the user in support of those purposes. Always, the overriding purpose is the answer to the question: "Why does the user want to create this visualisation?" Put another way, the same question can be stated: "What understanding of the dataspace is the user trying to achieve?"

2.1.2 Questions Based on the Outer Loop

In an evaluation using the VisTG Reference Model, each loop implies a particular set of questions the evaluator might ask, as follows:

(Note that the wordings here are not identical to those of the same questions in Chapter 1. The intent of the questions is the same, and the reader may derive that intent more precisely by comparing the different wordings than from either wording alone.)

- Q1. What user purpose is being considered?
  - Q1a. What higher-level purpose does this one support?
- Q2. What information does the user need to get from the computer to achieve the purpose?
- Q3. What does the user need to tell the computer to allow it to provide the needed information?
- Q4. What impediments might inhibit the user from taking advantage of the information provided?
- Q5. What impediments might inhibit the user from providing the computer the information it needs?

- Q6. Is there any mechanism to alert the user to information that might be important for the purpose, but that is not currently evident in the display?

2.1.2.1 Q1. User Purpose

The key to the entire evaluation procedure is to start with a defined purpose. If the purpose is insecurely defined, it becomes hard to determine what information the user needs in order to achieve it, and if the desired information is too vague, it is hard to know what the computer would need to be told so that it could provide the desired information.

To define the purpose is not necessarily as rigid as it may at first sound. A purpose such as “To discover possible causal relationships among events of types X and Y” may well be a purpose that can lead to a useful evaluation of some system. It specifies that the user needs to “discover”, which asserts that a wide range of data will probably need to be displayed; it mentions “relationships” which argues for the displays to be appropriate for coordinated presentation of disparate entities; and it specified “events of types X and Y”, which restricts the kinds of data required to those events and possibly the contexts surrounding them.

The foregoing purpose might be catalogued under “Research”, as it deals with general conditions that might at some time prove useful in an operational setting. A purpose with a more operational tone might be “to determine whether aircraft X has hostile intent” or “to know when an attack upon this network is being prepared”. The latter purpose suggests a second one: “to determine the nature of the attack upon this network that is being prepared”, which in its turn suggests a third: “to prepare defences against the attack being prepared.” Of course, for “network” one could substitute many other words, such as “position” on a battlefield, “theory” in an adversarial debate, “transportation hub” in case of a suspected terrorist attack.

This example illustrates one problem facing an evaluator using the VisTG Reference Model. Why is “to know when an attack upon this network is being prepared” a valid purpose that suggests further purposes, whereas “to know when an attack upon this network is being prepared and to determine its nature and prepare an appropriate defence” is not? This problem is not easily resolved. One clue might be the use of the word “and”, which in itself suggests a parallelism between the two parts of the purported purpose. Seeing the word “and”, the evaluator might well consider whether there were not two parallel purposes that might support one another, might conflict with one another, or might be independent but simultaneously support some higher purpose that ought to be defined. Another clue might come in the second stage of the evaluation, when the evaluator is answering question 2.

For a system to be designed, the purposes are specified by some technical authority, but in evaluating existing systems, they are specified by users. Typically, an evaluator might observe users, enquiring by minimally intrusive means as to what they were trying to achieve at any given moment. This is not an easy thing to do, since overtly asking a user, particularly a highly skilled user, is likely to distract from the user’s performance of the task. Highly skilled users very often are no longer conscious of what they are actually doing in support of their main purpose. Particularly in systems designed to support visualisation, they “see intuitively” the patterns that they can integrate into situation awareness, and have forgotten just how they came to have this ability. They may know that they are looking to see whether an incoming aircraft appears to have hostile intent, but not much more than that.

The problem facing an evaluator using the VisTG Reference Model and trying to define the actual purposes of a user is much the same as that facing anyone doing a task analysis. The evaluator must try to get at what is in the user’s mind while not affecting the very thing being studied. Different circumstances demand different tricks, but some interference is almost inevitable in the real world. It is a measure of the evaluator's skill that the purposes developed during the evaluation correspond to those actually held by the user.

One of the questions the analyst should ask about the purpose is why it exists (Question 1a). All purposes support some other purpose, even if it is as mundane as "to satisfy a superior officer". Nevertheless, in a complex system, the analyst may start by analysing a readily defined purpose, and then asking what other purposes this one supports and what other purposes support this one (the "why" and the "how" of this purpose). The set of purposive connections so developed eventually builds up a network of purposes that describes the entire operation of the system from the user's viewpoint. If that network turns out to be disjoint, the complex system might as well be two or more simpler systems. Each can be analysed separately. Usually, though, the "why" and the "how" develop into a single network.

Below, in reference to Q6, the same set of purposes is shown to have other possible mutual connections, of interference rather than of support. Those connections will form some kind of network, and the two together may form a foundation for the global analysis of the system. To date, no analysis has followed this route, at least in part because it requires automated support that does not yet exist.

2.1.2.2 Q2. Information Required for the Purpose

Having determined on a purpose in respect of which the system is to be evaluated, the evaluator needs to determine what the user needs to know in order to achieve it. This is the second of the standard questions: "Q2. What information does the user need to get from the computer to achieve the purpose?"

Of course, any user will have a lot of relevant information already in memory. The difference between a skilled user and a novice is how much is in memory and how much must be supplied by the computer. Inherent in this comment is that the evaluation intrinsically suggests requirements for user training or for design features that allow for differential interaction between user and computer according to the user's level of expertise.

Suppose the purpose is "to determine whether an approaching aircraft has hostile intent". The system evaluator must determine what information would enable the operator to make such an assessment. For example, a friendly plane may be moving according to a predetermined plan. Is the possibility of such a plan in the computer's dataspace, and is there a means whereby the operator could be made aware of it? A friendly plane may be carrying a responder that would send an automatic identification signal. Would the computer have the data as to whether such a signal was requested and if so, whether it was observed? Might the aircraft, if not friendly, be at least non-hostile, as, for example, a passenger airliner usually is? What significant clues might enable that distinction to be made? Is that information available in the dataspace, and if it is, can it be provided to the operator? Even if the aircraft itself is non-hostile, its flight pattern might suggest otherwise (as was the case in the attacks on the World Trade Center and the Pentagon in 2001). Is information about the aircraft's flight pattern in the dataspace, and can it be made available to the operator?

These are simply imagined examples that a skilled operator might be able to augment many times over. To determine what information an operator would really find useful for any particular purpose, the evaluator would have to observe or to ask. "Asking" often implies subtle experimentation and inference, rather than simply posing the question to the operator. People tend to ignore the obvious when asked such questions, since they "go without saying". If the answers are taken literally, some of the most important things might be omitted from the evaluation. Not only that, but when a person follows one line of thought, things that lie on another line may well be forgotten. Human Factors researchers have developed several techniques to surmount these problems, and they need concern us no further.

2.1.2.3 Q3. User Provision of Information to the Computer

It is all very well for the evaluator to know the user's purpose and the information the user would need to achieve it, but if the user cannot inform the computer that this information is required at a particular moment, there is little value in knowing that it exists. What information does the user need to provide to the computer that would let the computer determine that certain information should be provided to the user?

As with the previous question, the answer to this one depends a great deal on what information is already in the computer. A special-purpose system may only be able to produce one kind of information, and may have a display dedicated to presenting that information. Such a system is the limiting case. All the user needs to provide to the system is the information that the user is present, by switching it on. At the other extreme, a general-purpose data-handling system needs to be told in some detail the characteristics that would describe what the user wants to know. In the example of the aircraft, the user might need to let the computer know that a display of any interrogation and response sequences, with timings, was required if such was available. The user would presumably have to identify to the computer which of many potentially interesting aircraft was the one for which information was requested. Or, the computer might have been given prior knowledge, enabling it to know that the user ordinarily wanted to be informed about any aircraft for which the identification response failed to match a pre-specified pattern. In that case, the user would not have to inform the computer at the moment when the interrogation and response sequence was of interest for the user's visualisation of the situation.

The evaluator must consider, for each item of information required by the user, what information the user must give the system in order to get back the desired information. Note that at this level, the evaluator is not concerned with how the user provides the computer the information, or, for that matter, how the computer displays to the user its information.

The kind of information that the user must give to the computer in respect of a given purpose, if it is not already stored, is the same kind that would be given to a human. It should be at the same level of abstraction as the purpose itself. If the purpose is to determine the intent of an incoming aircraft, the kind of information required is relevant to the intent of the aircraft, not to the layout of a display of the information, and not to the organization of the dataspace. Those issues arise in conjunction with the analysis of the system's suitability for supporting those purposes of the user. Of course, those purposes arise out of the original "aircraft intent" purpose. In other language, they become sub-goals of the main goal of determining the aircraft intent.

Separating out the levels of purpose and supporting purpose, and keeping distinct the separate information and action requirements of each purpose and sub-purpose, turned out to be the most difficult aspect of using the VisTG Reference Model in practical evaluations of existing systems. There is a natural tendency for the analyst to think simultaneously of the information requirement and of the method by which that information is transmitted. The method is not relevant to the feedback loop's analysis sequence: Purpose – information the user needs from dataspace – information user must give computer.

2.1.2.4 Q4. Impediments to the User Getting the Information from the Computer

The simple loops of the canonical diagram of the VisTG Reference Model (Figure 2-1) camouflage much possible complexity in the real situation. One of the complexities hidden by the diagram is that either of the paths in each of the three main feedback loops could be compromised. The analysis of purpose may say that the user needs certain information, but the analysis of the real or future system suggests that the user may have trouble getting that information. In the "aircraft intent" example, the user may require information about the timeline of the aircraft's flight, but the system may not keep the path time-tagged (the information is not in the dataspace), the Engines in the computer may not be capable of translating the time-tagged path data into displayable form (failure of the information channel), or external interference may limit the reliability of the information (the information about the desired aircraft is mixed in with information about other aircraft in the channel for display, leading to clutter and "information overload").

The examples above are somewhat artificial. In particular, external interference is not very evident at this high-level of analysis. It applies much more when the analyst comes to the middle and inner loops (the Engine loop and the I/O loop). Then, the issue of information transmission may hinge on whether the user is in a noisy environment, whether the display colours are well chosen, whether the display device can be read under full sunlight, and so forth – but keeping the analysis to one level at a time, these factors merely suggest to the analyst that there may be problems in transmitting the information from the computer to the user. The analyst must then determine how those problems affect the user's ability to achieve the purpose being analyzed.

2.1.2.5 Q5. Impediments in the User Informing the Computer as to the Information Wanted

The user "talks to" the computer through input devices such as keyboards, pointers, body-movement sensors, microphones, and so forth. It is through these devices that the user informs the computer what information is wanted "right now". It is tempting for the analyst to think first about the devices when considering the information the user needs to transmit, because it is usually the devices that limit the user's ability. Confusing the device limitations with the information requirements proved to be one source of difficulty in using the VisTG Reference Model for practical evaluations.

In analyzing the loop associated with one purpose, only information relevant to that purpose should be considered. If the user wants to see a flight path time-line, the analyst must determine whether there are impediments to the user making that known to the computer, not whether or how the user can encode that instruction – not whether there is an appropriate menu item, or whether the computer is able to decode spoken language. Those may in fact be the reasons that the user has difficulty making the information requirements known to the computer, and they will be adduced at a later stage of the analysis, but at this level, the analyst's question is how the impediments affect the user's ability to get the required information and thereby affect the user's ability to achieve the purpose.

There is a clear relationship between this question and the preceding one. If, for example, the user wants to see a time-tagged flight path for a particular aircraft, and has no means to specify the aircraft or to specify that the flight path should be displayed as a time line, then the computer will not provide the information the user needs. That is not an impediment on the path through which the computer provides the information, even though the result is the same as if there were an impediment on that path. Making this distinction between an inability of the computer to provide information and an inability of the system to receive requests for information is another area that may cause difficulty to the analyst.

2.1.2.6 Q6. Alerting the User to Possibly Useful Information

The preceding five questions relate to the loop supporting a single purpose. Usually, however, the user is pursuing more than one purpose at a time, though not paying attention to more than one of them (as shown in Chapter 1, Figure 1-7). Even if there is only one overt purpose, in many cases the user cannot see all of the relevant information in a single display, or there may be so much data on the display that significant patterns are likely to be overlooked. Any of these conditions mean that information relevant to a purpose may not get into the user's mental picture. To cover these possibilities, many systems offer alerting mechanisms. These may be a simple as the visual or acoustic signal on a home computer that indicates new mail has arrived. The signal draws attention, and the user has the opportunity to drop what she is doing and glance at the mailbox to see whether she wants to pay attention to the newly arrived message.

Alerting systems have two common characteristics:

- The system's Engines act autonomously once the criteria for the existence of an alert condition have been specified and the Engine activated; and
- The alert indication intrudes upon the user's attention and offers the user an opportunity to determine whether the event that caused the alert warrants a further diversion of attention from the ongoing task.

Other than these two characteristics, alerting systems come in many different forms. Biological alerting systems in the human, for example, include built-in alerters such as the detection of unexpected motion in the visual periphery, or the cessation or onset of a continuing sound. In addition, a person can set an alerting condition. One hears the phone ring while hosting a party much better when one is awaiting a phone call than if it comes out of the blue. Computer-based alerting systems likewise may be built-in (programmed by the developer) or temporary (specified by the user).

Setting an alert condition is a purpose on its own that the analysis must at some point consider. It is not, however, part of the purpose for which the alert may be useful. Continuing the example of the intent of an incoming aircraft, a suitable alert might be one that is set by the computer when it autonomously finds in the dataspace some data that fit a criterion that applies to aircraft that have a significant likelihood of hostile intent. Those criteria might have been set recently by the current user, or they might have been a consequence of long development by intelligence officers before the system was even built. Either way, they serve the purpose of determining the intent of a potentially hostile aircraft, by letting the user know that this purpose might reasonably be reactivated at a time the user was doing something else (perhaps even evaluating the intent of a different aircraft).

If an alert for this purpose can attract the attention of a user performing a task in support of a different purpose, then alerts on behalf of other purposes can likewise interrupt the performance of tasks in support of this purpose. They form potential interferences to the transmission of information from the computer (Q4 in the canonical list of questions the analyst should ask). Furthermore, any actions the user needs to take in order to determine whether the alert is worth pursuing may interfere with actions in support of the current purpose (Q5 in the canonical list). Alerting systems therefore form a network of linkages among the different purposes.

It is a commonplace that excessive alerts tend to be ignored, especially if they are too intrusive, or if dismissing the uninteresting ones takes too much effort away from an ongoing task. It even has a centuries-old name: "Crying Wolf." Pilots are well known to turn off intrusive acoustic alerts that might save them from a crash, but which occur in many non-dangerous or deliberately caused situations. Nevertheless, a real wolf might be eating the sheep; the plane might be closer to the ground than the pilot thinks. The system analyst's job is to determine the degree to which the occurrence of possible alerts is likely to detract from the user's perception of information about the current purpose, the degree to which actions involved in responding to the alert interfere with actions in support of the current purpose, and the relative importance to the user of the current purpose and the purpose on behalf of which the alert is provided.

When considering alerts, the relative importance of the different purposes is something that often can be determined only from real users. A system might allow users to turn off alerts for purposes sufficiently less important than the current one, while allowing those same alerts to occur when the user switches attention to a less important current purpose. The degree to which this happens will depend on several factors, including at least: the probability that an alert really does represent the condition for which it was designed; the interference with the user's attention caused by the alert; and the interference with the user's actions caused by responding to the alert sufficiently to see whether it truly represents a condition that justifies switching attention to the purpose on behalf of which the alert was issued.

2.1.2.7 Canonical Questions Summary

An analysis using the VisTG Reference Model starts by selecting a purpose. Purposes are achieved by acquiring information from the computer, and information is acquired from the computer by telling it what information is required. Both processes are subject to possible impediments, either internal to the transmission process or as noise from outside sources. The analyst must consider all elements of this loop, to determine how readily the user can achieve the purpose.

Each purpose supports at least one higher-level purpose (why the purpose exists), and except at the level of physical interaction with the computer, each is supported by at least one lower-level purpose (how the purpose is implemented). The analyst eventually must trace the network implied by these connections of support, to build a picture of the entire system as seen by the user. However, in analysing the suitability of the system for a particular purpose, the up-going "why" connection may be safely ignored.

A purpose may require the user to examine more data than can readily be displayed or attended at any one time. If the computer can be told what characteristics are likely to signify portions of the dataspace likely to be useful for the purpose, it might issue an alert when those characteristics appear in a portion of the space seeming not currently to be in the user's attention. Since an alert is designed to attract the user's attention away from the current task, alerting possibilities create a second network, of inhibitory interconnections, among the purposes supported by a complex system. This network must be considered when the analyst deals with the possible interferences in the primary loop for any specific purpose being analysed.

2.1.3 Levels of Analysis

The VisTG Reference Model diagram shows three levels of feedback loops, an outer one between human understanding and the dataspace, a middle one between human visualisation and the Engines, and an inner one concerning the I/O devices through which all information is ultimately passed between human and computer. This three-level structure simplifies the representation of the model, but has been shown to be tricky to use in practice. The Norwegian delegate attempted to use the VisTG Reference Model in evaluating the NORCCIS II system, and commented (see Annex):

> I was not able to "evaluate" NORCCIS II according to the VisTG Reference Model and the related procedures we have. I really lost the battle when thinking about how to do it in a meaningful way. "[The] 6 questions" immediately produced a lot of answers, but I could not organize the lot of them in the structure of the three level model.

It may have been a mistake to treat the simplified three level structure initially as if those were the only levels of analysis. The intent was to separate the interactions that (1) involved the data and the human's understanding of the data (outer loop) from (2) the interactions that involved the computer's manipulations of the data and the human's visualisation using the manipulated data (middle loop), and both from (3) the interactions through the physical interface devices (I/O loop). Each of these "levels" might contain several levels of supporting and supported purposes and the feedback loops inherent in each purpose. Figure 2-2 shows a simple two-level structure in which three lower level purposes are defined by the action set that supports a single higher-level purpose.

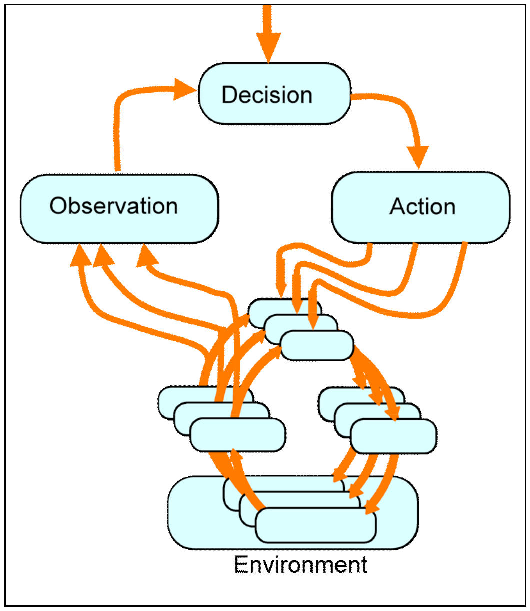

Figure 2-2: Multiple Purposes Supporting One Higher-Level Purpose. The higher-level purpose is symbolized by the topmost downward-pointing arrow and the supporting purposes by the three arrows leading from the higher-level action.

2.1.3.1 Supporting and Supported Purposes

In the analysis of the feedback loop associated with a single purpose, through the six questions, the analyst will have come across several instances in which certain items are required. Usually, these take the form of actions on the part of the user. Continuing the example of identifying the intent of an aircraft, one of the supposed requirements was that the user tell the computer which aircraft was under consideration. This requirement identifies a new purpose that supports the original one. A new purpose implies a new iteration of the six questions, which might be paraphrased and answered as follows:

- Q2.1. What is the purpose?

  A2.1. To ensure that the computer reports information on the correct aircraft.

  - Q2.1a. *What purpose does this support?*
  - A2.1a. To determine whether a particular aircraft has a hostile intent.

- Q2.2. *What information does the user need?*

  A2.2. To see which aircraft the computer is referencing and describing.

- Q2.3. *What does the user need to tell the computer?*

  A2.3. Which of the possible aircraft is the one of interest.

- Q2.4. *What impediments might inhibit the user determining which aircraft is being described?*

  A2.4. The interesting aircraft may not be presently displayed; the display may not differentiate the aircraft about which the computer is providing information from others being displayed; if the aircraft is being displayed and is distinguished, there may be too much else on the screen for the user reliably to discern which one is being singled out.

- Q2.5. *What impediments might inhibit the user telling the computer which aircraft is of interest?*

  A2.5. The display of the desired aircraft may inhibit selection by pointing, because of resolution or ambiguity issues; if the desired aircraft is not being displayed, the user may have no mechanism for identifying it; if identification is by pointing, the user may need to change the geographic region of the display in order to bring the desired aircraft into view; if the interest of the aircraft has been proposed by the computer through an alert, the user may have no current means of accepting or rejecting the alert, or of choosing to examine it further, because the information channels are being used for other purposes.

- Q2.6. *What alerts are provided?*

  A2.6. Perhaps the system has thresholds set for flight envelopes associate with potential hostile intent, or for failure to provide proper responses to interrogative signals, or has sophisticated algorithms for assessing the likelihood of hostile intent. In some manner, it signals the user that its examination of the dataspace suggests that there exists a potentially hostile aircraft not currently the focus of the user's attention nor previously dismissed as non-hostile on the basis of the same evidence. An alert signal would be presented when the user is working with the feedback loop supporting some purpose other than ones that support the "top-level" purpose of determining the intent of the aircraft.

  By identifying the possible existence of such an alert, the analyst has identified the possibility of interference that must be taken into account when quite different purposes are being analyzed. The proliferation of such possibilities into a substantial network is one area in which automated support is almost essential.

2.1.3.2 Interference Among Loops

The necessity for considering alerting signals points up that the user is unlikely to be able to see all the relevant data at once. Usually, several different things are going on. The user may well be pursuing several simultaneous purposes in support of the main task, as suggested in Figure 2-2. Attention to one purpose will usually detract from attention available for another. The very existence of an alerting signal is designed to draw attention away from whatever purpose the user is currently pursuing, and must reduce the user's effectiveness in that pursuit. The analyst must consider the possible interactions among the different feedback loops caused by the virtual network of alerting interferences.

Alerts and the division of attention are not the only source of mutual interference among the various purpose-drive feedback loops. Different loops may need to share, at some point, information channels of limited capacity. This kind of mutual interference happens largely at the level of physical input-output devices. Screen real-estate may be limited, and the user has only two hands with which to control a keyboard, a mouse, and perhaps other devices. These devices may be needed for communication to two or more simultaneously active algorithmic Engines, but at any moment can be communicating only with one. Even if several devices are available, the user would have to spend time switching from one to another, not to mention the mental load that would be incurred in doing so.

Interference among loops is something that must be considered in answering canonical questions 4 and 5, about impediments to effective information transmission, but in the initial stages of an analysis, when only a few purposes have been analyzed, the possibilities for mutual interference may not be readily observed. Again, automated support would be an advantage, using the fact that at some point in the analysis, all the transmission paths must be identified, and an automated system could, in principle, detect when two or more purpose-support structures use a common path.

2.1.3.3 Interactions within the "Engines" Loop

The examples in the previous sections dealt briefly with a purpose that supported a higher level purpose, but both were within the loop concerned with the interactions of the user with the database. In other words, they belonged within the outer loop of the VisTG Reference Model as shown in Figure 2-1. Some of the answers to the questions did, however, suggest the existence of purposes, and therefore of feedback loops, within the middle loop of Figure 2-1.

Consider, for example, the suggested answer to Question 2.5 above, which includes: "*if identification is by pointing, the user may need to change the geographic region of the display in order to bring the desired aircraft into view*." In changing the geographic region, the user is not interacting with the dataspace content, but with the selection of data from the dataspace and its availability for display. Selection and manipulation of the data are the functions of the Engines, and the user may have a hierarchy of purposes concerned with those functions.

Following through the same example, the answers to questions 2.1 to 2.6 suggest new purposes, some of which relate to Engine function – in other words not to understanding the dataspace, but to organizing it for better understanding. There are several kinds of Engines, but most non-trivial systems include at least:

- Navigation Engine: Navigate to the relevant part of the dataspace;
- Selection Engine 1: Select the relevant data;
- Selection Engine 2: Select algorithms to apply to the selected data; and
- Various Engines: Execute the appropriate algorithms on the selected data.

In the example, the cited statement suggests that the user will be concerned with a Navigation Engine, one that is able to "change the geographic region of the display". Accordingly, a new set of six questions automatically arises.

- Q3.1 *What is the user's purpose?*

  A3.1 To move the geographic region of the display so that it includes the aircraft of interest.

- Q3.2 *What information does the user need?*

  A3.2 The geographic location of the aircraft, either in absolute terms or relative to the geographic region currently displayed.

- Q3.3 *What does the user need to be able to tell the computer?*

  A3.3 The user needs to be able to tell the computer in what way to change the area being displayed, whether to a region specified by geographic coordinates, to a region specified by its relation to the currently displayed region (zoom and pan, colloquially), or to a region specified by the computer in relation to an aircraft identifiable by means unrelated to its geographic location.

- Q3.4 *What impediments might inhibit the user's ability to see what is necessary?*

  A3.4 The user might not be able to see where the aircraft is relative to the currently displayed region.

- Q3.5 *What impediments might inhibit the user's ability to tell the computer what it needs?*

  A3.5 The user might not have a means to specify a region to display relative to an object in the region.

- Q3.6 *What alerts might aid the user's purpose?*

  A3.6 Alerts most commonly occur in response to the content of the dataspace, and are therefore not usually associated with Engine operations. There may be Engine conditions that suggest the occurrence of an alert, as, for example might be the case with an algorithmic Engine whose operations continue in the background until completion, in which case the user might well be alerted that the operation is complete – but for the purpose under consideration here, and for most Navigational Engines, there seems little occasion for alerting.

Once the user has used the Navigation Engine to bring the right part of the dataspace into the display, other Engines will probably prove to be useful in support of purposes adduced in the earlier parts of the analysis. For example, a Selection Engine is needed if the user wants to select data relevant to a particular aircraft. Algorithm Selection Engines may be involved in choosing ways of analyzing the data from the aircraft, and algorithmic Engines in manipulating the data in conjunction with the user's own experience in order to make the final assessment of the aircraft's intent.

Much of the user's interaction will be with the algorithmic Engines, and to a large extent in some systems, it will be through the algorithmic Engines that the user interacts with the dataspace. The algorithmic Engines may (and probably will) themselves use Selection Engines without overt user interaction, but these interactions concern the programmer more than they concern the analyst using the VisTG Reference Model in considering the value of a system for a purpose. What that analyst wants to know is how well the algorithmic Engines do what the user believes they are supposed to do.

2.1.4 The Input-Output Loop

All the information transmission between human and computer mentioned in the foregoing must pass through physical devices. The algorithmic Engines are responsible for selecting data from the dataspace and doing the computation necessary to bring the data into a form that matches the user's purpose. The job of the Input-Output loop is to take the processed data and present it on the output devices, or to accept the user's input and turn it into a form suitable for the algorithmic Engines to use. The Input-Output loop is therefore more than just the connections of the physical devices. It concerns, for example, the layout of presentations on a screen, including the simultaneous presentation of data from different Engines, as for example in Figure 2-3. Much of the bandwidth limitation that induces interference among the loops supporting different purposes occurs at this level.

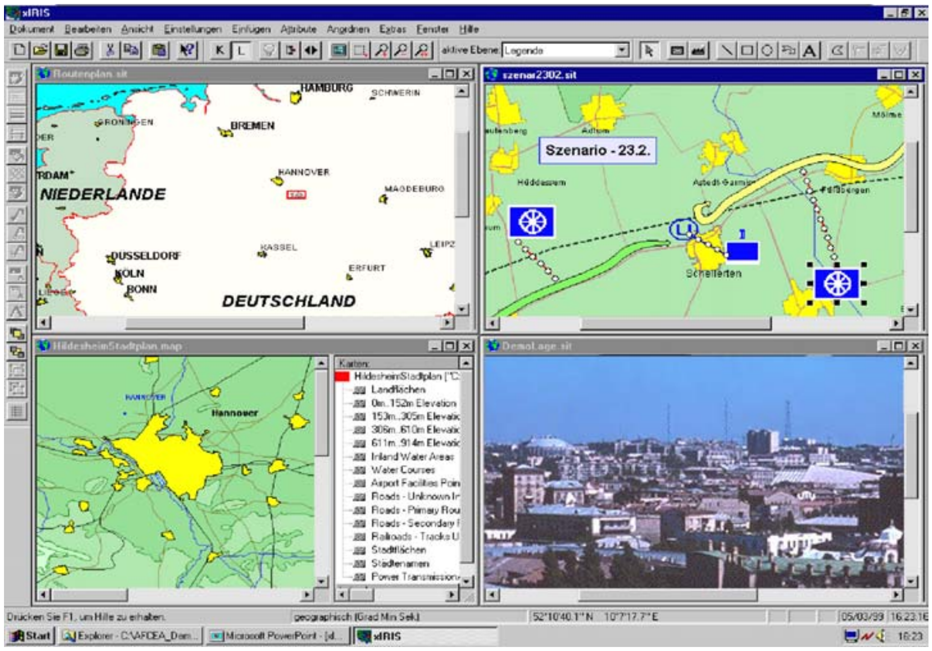

Figure 2-3: A Screen from the German xIRIS Project Showing Simultaneous Output from Different "Engines" in the Computer (FGAN-FKIE, reproduced from the Final Report of IST-013/RTG-002, published by RTO as RTO-TR-030).

"Visualisation" is often taken to mean "presentation" on a display. In the VisTG Reference Model, this is one side of the innermost loop of the structure that supports the visualisation that the human does in the head. Even though it is but one facet, however, it is an important one, as it is the channel whereby the data extracted from the dataspace, and manipulated by the Engines, is provided to the human. If that transformation, from machine representation to human perception, is ineffective, no amount of cleverness in the rest of the structure will compensate.

The Final Report of IST-013/RTG-007 (colloquially known as "the HAT Report", or formally as RTO-TR-IST-030, "Visualisation of massive Military Datasets: Human Factors, Applications, and Technologies") presented a taxonomy of data, display, and presentation types that was intended as a basis for developing presentations suited to different tasks of the user and to different display technologies. Here, we present a brief summary and perhaps a slight update to that taxonomy.

2.1.4.1 Displays and Presentations

The word "display" carries at least two very different connotations: a piece of hardware that produces pictures and a particular presentation shown by that hardware. In this section, "display" refers to the hardware, and what it shows is a "presentation" that allows a user to visualise something about the task at hand. A presentation uses a display to allow the user to visualise the implications of data.

Data have many forms and interactions with the user and with the real world they represent. The Final Report of IST-013/RTG-007 (Visualisation of Massive Military Datasets: Human Factors, Applications and Technologies, RTO-TR-030; colloquially known as "The HAT Report"), categorized them on six dimensions, all of which have implications for displays and presentations.

These dimensions affect considerably the kind of presentation, the kind of display, and the kind of user interaction that is most effective for a particular task. For a full discussion, see the HAT Report, Chapter 3. Here, it may suffice to provide a brief description.

- **Acquisition**: Streamed data are being acquired while the user's task is in progress. Static data are assumed to be invariant for the duration of the task.
- **Sources**: The source of data may be single (meaning that all data elements are of the same nature) or multiple (meaning that the data are heterogeneous, either in type or in incorporating the source identification within the data element). If the Acquisition parameter is "streamed" and the Source is "multiple", then a sub-categorization may be warranted: whether the data from different sources are regularly or irregularly placed with relation to each other.
- **Choice**: Does the user choose which data are to be displayed or is the choice imposed on the user? Strictly speaking, this may not be a characteristic of the data, as in a briefing the briefing officer may select the displayed data (which makes the data user-selected), while the audience sees what is provided for them (making the data externally imposed). The argument for asserting "Choice" to be an attribute of the data is that the task of the briefing officer differs from the task of the audience, and the "Choice" mode of the data affects each kind of user's appreciation of what the data imply.
- **Identification**: The data may be identified by a label (such as "City: Prague") or it may be identified by its location in some underlying continuum (such as 0300 hours). A datum may have both characteristics (such as City: Greenwich: Longitude 0 degrees).
- **Values**: Data elements may be structured, and usually are, as in the previous example. However, the atoms of the structure are either analogue (values in a continuum of possibly many dimensions) or categorical (having membership in a definable class). Technically, these are not mutually exclusive, as 1.05 is a member of the class of numbers less than two, but in most cases of interest for visualisation, an atom of a data structure is a representation either of a point in a continuum or a selection of one out of a finite number of possible categories. If the Value is Categoric, it may be symbolic, the category representing something else in the way that a word is not the thing it represents, or it may be non-symbolic (though its representation on a display must intrinsically be symbolic).
- **Interrelations**: Data elements seldom are of interest individually. Their context usually is what gives them meaning. The structure of the relations among the elements may be developed by the user in the process of analysis, or it may be intrinsic in the source. For example, in a physical network, two nodes are or are not connected by a single link, whether the user likes it or not, but in a conceptual network being developed by a social analyst, the structure may be largely created by the analyst.

Table 2-1: Dimensions of Data Description

Acquisition	Streamed	Regular	
		Sporadic	
	Static		
Sources	Single		
	Multiple		
Choice	User-selected		
	Externally imposed		
Identification	Located		
	Labelled		
Values	Analogue	Scalar	
		Vector	
	Categoric (Classical or Fuzzy)	Symbolic	Linguistic
			Non-linguistic
		Non-symbolic	Linguistic
			Non-linguistic
Interrelations	User-structured		
	Source-structured		

2.1.4.2 Types of Display

Presentations also vary on several dimensions of description, some of which is essentially parallel to the dimensions of data description. The HAT Report listed:

Table 2-2: Presentation Types

Display timing	Static	
	Dynamic	
Data selection	User-selected	
	Algorithmically directed	
Data placement	Located	
	Labelled	
Data values	Analogue	Scalar
		Vector
	Categoric	Linguistic
		Non-linguistic

This categorization focused on the relationship between the user, the data, and the presentation. It did not consider different kinds of hardware that might implement the requirement for representing, say, user-selected, dynamic, located, analogue data with a categoric overlay (such as an interactive map display with a symbolic overlay).

The hardware of displays covers a wide range, from small hand-held devices to wall-sized screens to "sandbox" 3D displays to immersive "caves". If one includes the supporting software, they also include variations in the degree of control the user has over the display (as opposed to the control the user has over the data displayed). Considering the different possibilities, a second table can represent display hardware:

Table 2-3: Some Characteristics of Displays

Display size	Hand-held
	Desk monitor
	Table-top
	Wall or Cave
Dimensionality	2D
	3D
User involvement	Exocentric view
	Immersive

Of course, not all combinations of these characteristics make sense – it is hard to imagine a hand-held 3D immersive display, for example – but most combinations do refer to feasible or existing displays. The larger the display, the more reasonable it is to consider an immersive system (one in which the user feels part of the scene). To judge from the video game industry, even desktop displays can become quite immersive. The distinction between exocentric and immersive is not strictly a characteristic of the display hardware, but the hardware may determine whether immersivity is easily achieved. Immersivity exists when the user has the feeling of being in the scene being displayed. A 3D display need not be immersive. A "Cave" usually is, since in the Cave a fully 3D display surrounds the user on all sides (except directly behind), but a "3D sandbox" usually is not. Users of a 3D sandbox look down upon or into a volume within which the display exists, taking a "theoscopic" or "God-like" view on what is being displayed. They may be able to act on entities within the space, but they do so in the manner of a person playing with toy soldiers, whereas in a Cave, the virtual soldiers may seem to exist within the same space as the user.

In a 2D display, all entities are formed on a single surface, which means that the user can often select the appropriate entity by simple pointing, though the result may sometimes be ambiguous: when one points at a location on a map, is one pointing at a location on a road, the road, the town in which the road lies, or the county in which the town lies? Such ambiguities, while possible, are not usual in references to entities in 2D displays. They are endemic in interactions with entities in 3D displays. There is no difficulty if a virtual object is displayed as if it were a discrete object in real space; it can be picked up as one might pick up a real object – but since the display is of a virtual volume, simple pointing is problematic. The direction of the point extends to infinity, and in a virtual space, the first surface encountered by the pointing line may well not be the surface of the intended object. Surfaces may be transparent or partially so, entities may be diffuse or even conceptual (for example, the virtual air in from of a simulated aircraft might contain a dense mesh of data points representing the probability of coming under ground fire, or something like simulated smoke might represent the density of toxic pollutants downstream of a source), or the volume may have density variations that represent aspects of the data being displayed. In none of these cases is it necessarily obvious how the user can interact with specified elements of the presentation.

Different display characteristics are appropriate for different display types. It is impossible, for example, to have an immersive display with static data, or with algorithmically selected data. For an immersive experience, the user must interact with the data at least insofar as to change the viewpoint in a reasonably natural way. On the other hand, many of the different combinations of data and display types work well together, as shown in Table 2-4, again from the HAT Report.

Table 2-4: Some Characteristic Relations Among Data and Presentation Types

Data Type	Presentation Type	Comment
Streamed	Dynamic	The user ordinarily wants to act when some event occurs.
Located 2D or 3D	Located	The display is a 2D or 3D map of some attribute(s) of the data. If the location identification of the data is in a higher dimensional space, there is no such natural mapping. Tricks must be used.
Labelled	Labelled	The display is likely to be tabular, or some kind of a graph such as a histogram or pie chart.
Analogue scalar	Analogue scalar	Even if the data are identified by label, its analogue values map naturally to analogue display variables such as the length of a line or the brightness of a pixel.
Analogue vector	If 2D or 3D, Analogue vector	A 2D attribute can be mapped onto an area display, a line with length and orientation, a colour hue, a sound location, a sound intensity and pitch, and so forth, all analogue vector attributes of the display. A 3D attribute can similarly be mapped into a volumetric display. Higher dimensional analogue attributes can be displayed, but the mapping is less obviously "natural."
Categoric	Categoric	Categoric data values have no natural relation to analogue display values, and must be displayed categorically. The categoric display attributes may or may not map "naturally" onto the categoric data attributes. This kind of mapping is usually considered to be "cognitive metaphor."

2.1.4.3 Fisheye Views

The term "Fisheye View" refers to a representation of a dataspace in which a small "focal" region is displayed in considerable detail, while a contextual region – possibly incorporating the whole dataspace – is simultaneously shown at progressively lower resolution as the distance from the focal region increases. It is not clear why the term "fisheye" has come to be associated with focus-plus-context displays, because a "fisheye lens" does not work this way, whereas our human eyes do. Our eyes have a very small central region that sees at high resolution (the fovea), surrounded by a wide region covering nearly a hemisphere at progressively lower resolution. Despite this, we do not usually notice that only a very small part of the world is seen at any moment at high resolution. Why not? What allows us to see our world as a high-resolution whole? Can we create displays that provide the user the same ability in a more abstract dataspace?

The human visual system has three important characteristics: the first is that the high resolution of the fovea is carried through the various stages of visual processing. The second is that in the low-resolution part of the retina, the processing system is arranged so that the locations of potentially interesting events are signalled. The third is that the eye is a lightweight sphere in a well-lubricated socket, with strong muscles that can move it quickly from one pointing direction to another. In conjunction, these characteristics mean that the eye can very quickly and easily be redirected so that the focal region is briefly aimed to see at high resolution whether a signalled event indicates that deeper examination might be useful, and equally quickly be returned to the original aim if the event turns out to be insignificant. The memory of the high-resolution glance in the shifted direction contributes to the perceived view of the space around us, at least for a short while.

It is this coupling between autonomous event processing and rapid, easy, redeployment of the focal area that makes our visual focus-plus-context representation useful. If the eye were heavier, requiring effort and the control of inertia to shift its direction quickly and accurately to a new focal point and back again, or if the muscles were weaker, or, most importantly, if there were no signalling mechanism in the low-resolution part of the visual field, our human “fisheye view” would be much less useful.

Interaction is inherent in the very idea of a fisheye view, even in views on more abstract dataspaces. The simultaneous display of the context and the focal region ordinarily implies that the user may want to change which area constitutes the focus. Often, that change needs to be rapid and effortless, with an equally easy reversion to the original focus location, as is the case with a flick of the eye. This implies that the user not only must be able to see in the context reasons why the focal region might need to be shifted, but also must be able to see how to set the focus accurately to the potentially important region and back to the original location. These requirements constrain how the context is displayed in any particular fisheye implementation.

Fisheye views may be implemented in many different ways. Here are a few real or hypothetical examples:

- A textbook might be displayed in full text for a few lines, surrounded by the subheadings in the same section, the main headings within the chapter, and the chapter headings for the whole book.
- Alternatively, the same textbook might be displayed with a central block of full text, surrounded by summaries of conceptually related material. The “fisheye” here would be in the space of concepts rather than in the space of literal text.
- An object-oriented software structure might be displayed as a graphical network showing all message and inheritance paths directly associated with a small chunk of textually displayed code, together with “trunk” paths linking the local areas with other blocks of objects, and those more distantly associated blocks with the operating environment of the software.
- A terrain map could be displayed at 1:1000 resolution in a central area, diminishing to 1:100,000 around the edges of the display. The popular “Falk Plan” maps of European cities often have a mild form of this kind of nonlinear magnification, showing the dense old core of the city at high resolution with smoothly reducing resolution for the outer and then the suburban regions.
- A sociogram could show the interactions of an individual with a few other individuals who form a close-knit group, of that group as a whole with other small groups that form a subculture, and of the larger group with other cultures and nations.

- A stock-market display could show detailed within-day trading data for one stock, with lower resolution data for the preceding week, and week-by-week data for the preceding year, while at the same time showing in a different dimension lower resolution trends for stocks of similar companies, and comparing those trends with data for other kinds of stocks at ever lower resolution depending on the "similarity distance" to the focal stock.
- A transportation network display may show detailed time schedules for connections between specified cities within a small time window, while showing less detail for connections nearby in time or to cities near the destination and for possible extensions to the trip.

What all these displays have in common is that they are most useful when the user has a special, though possibly momentary, interest in the focal region, while still needing to see aspects of its context either simultaneously or in the near future.

Why would a user want to seethe focal region in a low-resolution context rather than extending the focal area of fine detail over the whole display? A simplistic answer is "Data Clutter" sometimes called "Information Overload." One cannot deal with too much irrelevant detail. The irrelevant tends to obscure the relevant, or at least to demand effort in distinguishing which is which. No matter what the dataspace, the user is always dealing only with one aspect of the data at any one moment, though that aspect may be at a high level of abstraction. So, given that the whole dataspace usually cannot be shown in full detail (and should not, even were it possible), why is it better to show a decreasing-resolution context rather than a larger focal area at constant high resolution?

There are two classes of reason:

1) The wider context improves the ability of the user to evaluate the implications of the data in the focal region, and
2) The user may be interested not in that specific focal region, but in identifying in the dataspace those focal regions with characteristics that elsewhere we have labeled "Danger and Opportunity" (DAO).

Reason 1 applies most often when the user's interest in the focal region includes the relationship between its characteristics and the local variation of those characteristics. Reason 2 applies under many different circumstances, particularly if the user wants to look for specific information, to explore different areas of the dataspace, or to determine whether an alerting event is worth attention.

Conversely, why would a user not want to see a high-resolution central area in a lower-resolution context? A simplistic answer is "Structure distortion." No matter whether the "fisheye" is a nonlinear magnification of a geographic terrain or an abstract representation of some conceptual structure, the differential representation of data in different regions of the dataspace inevitably distorts something about the relationships among the regions. In terrestrial mapping, the common Mercator projection faithfully reproduces the orthogonal relationship between lines of latitude and longitude, which grossly distorting the areas of regions in different latitudes, whereas an equal area projection is likely to be cut into segments, or to distort the shapes of different regions. If what the user wants to know is inherent not in the content but in the structure of the data, a constant but low resolution display of a large part of the dataspace may be more effective than a fisheye representation that encompasses the whole dataspace.

Outside the computer application, the effect of narrowing the visible context can be seen in the difficulties helicopter pilots often have when using night-vision goggles that have a field of vision much narrower than the 210 degrees available in normal daylight vision. The focal area is unchanged, but the loss of the very low-resolution part of the peripheral context makes the pilot's task much more difficult. Similar difficulties may well occur when computerized displays show only a region of uniformly high detail, leaving the perception of the context to the user's memory or imagination.

2.1.4.3.1 Fisheye Versus Zoom

Under what circumstances is it better to display a fisheye view than to allow the user to zoom in and out of the dataspace, showing large parts of the space at low resolution at one moment and at the next a small part of the space in great detail? Can fisheye be combined with zoom?

What is important about the "fisheye view" is not the display itself, but the availability of information on which the user can base future action. We have argued that there are four different kinds of uses of information – perceptual modes: controlling/monitoring, searching, exploring, and alerting. The fisheye view supports them all, whereas a zooming display at fine detail supports mainly monitoring/controlling, and at low resolution supports mainly searching and exploring.

Alerting, as such, demands no specific support; what it does require is the ability for the user rapidly and easily to focus on the area indicated by the alert, and this involves a search (low resolution) and monitoring (high resolution) sequence of operations. In a zooming type of presentation, an alert relevant to an undisplayed region of the dataspace requires the user to zoom out, identify the region of the dataspace associated with the alert, move the target area to that location, and zoom in to it. In a fisheye representation, the user only needs to identify the region of the dataspace and move the focus of the fisheye there.

2.1.4.4 Network Representation

One of the data dimensions of Table 2-1 that is ignored in Table 2-4 is that of interrelation. Interrelationships are of many kinds, and most of the problems of effective display are concerned with how one aspect of the data relates to another. It may be the relationship of a graphic view of a scene to its neighbouring geographic context; it may be the relationship of a data value to its reliability or to its uncertainty (or to both); it may be the relationship of a condition to the probability of a future event; or it may be a different kind of relationship altogether. Not only the fact of a relationship may need to be displayed, but its nature may not be implicit in the way that the relationship is between a graphic scene and its surrounding geographic context. How, for example, should one consider or display the relationship between the commander of a peacekeeping unit and the mayor of a town in conflict?

A relationship exists between two data elements, but another way to look at it is to say that the two elements are linked in a definable way. The two elements can be described as nodes, and the relationship as a link. If one or both of the nodes has a relationship with some other element, the resulting complex is a network.

Networks have many forms, differentiated by the properties of their nodes and links. A node may be atomic or structured, and if structured, its atoms may be any of the types shown in Table 2-1 ("Interrelations" obviously referring to the structure, not to the atoms individually). Links between nodes also have properties, properties that depend to some extent on the properties of the atoms they relate.

Working Group 5 of the Workshop IST-043 organised by IST-021 identified the set of properties identified in Table 2-5 that might be attributed to links in a network, along with an indication of subsets of those variables that might be important in considering how to present visual representations of networks in three different application areas (Counter-Terrorism – CT, Information Assurance – IA, and Logistic Support Management – LS):

Table 2-5: Some Variables Associated with Networks of Relationships

Constraints	Rules			
Nodes	Location	CT		LS
	Node Type			LS
	Symbology			
	Open/Closed		IA	
	Node ID	CT		
	Input/Output property	CT	IA	LS
	Number/Type of links	CT		LS
Links	Capacity			LS
	Weight			LS
	Strength	CT		LS
	Direction	CT		LS
	Availability	CT	IA	LS
	Type of Traffic	CT		LS
	Location			LS
	Route			LS
	Identification (unique)			
	The Medium	CT		
Thresholds and changes		CT		
Traffic	Flow		IA	LS
	Path	CT		LS
	Routing			LS
Topology	Symmetric/Asymmetric	CT		
	Boundaries			
	Layout	CT	IA	LS
	Hierarchic/Non-Hierarchic	CT		
	Tree		IA	
	Topology Evolution in Time and Space	CT	IA	
	Partially Connected			
Logical/Physical				
Redundancy			IA	LS
Protocols		CT	IA	
Network Interconnections		CT	IA	LS
Location		CT		
Open/Closed			IA	
Layers				
Hierarchy		CT	IA	

One aspect of networks not reflected in Table 2-5 is their extension in time. When one considers time, several other factor or properties come into play. The network may be defined by a relatively fixed structure such as a set of wires. The representation of such a network through time would show the entire set of wires, with occasional moments at which a wire was added or deleted from the network.

On the other hand, a network such as a social network is defined by contacts among people. These are effectively isolated events. A talks to B, and on another occasion A may talk to C. An ordinary sociogram that ignores time would show A connected to B and to C, but the interpretation of the events might critically depend on which conversation occurred first. The network is ephemeral. A may have a communication channel with B that is continuously available, as is a wire, but the network of actual communication is a distinct network. This difference is critical in the detection of attacks in cyberspace. The network of connections between end-user computers through various routing nodes can be mapped, and that mapping might be useful for some purposes – but detection of an attack depends critically on the temporal sequence of packets traversing this network. The traffic itself constitutes a network.

Stigmergic systems provide another type of network. The prototypical instance is a termite mound, created by a process in which a termite that deposits a mud ball leaves a scent that increases the probability that another termite will drop its mud ball on top of the first, or, if on top of a nearby one, will drop the new mud-ball in such a way that the column so formed leans toward the other mud-ball. The stigmergic field is the region affected by the scent of the first mud-ball, and it affects termites not specified by the action of the first termite. A broadcaster does not specify the recipients of a broadcast message, but the broadcast event does affect the recipients at the time of the broadcast. Action in a stigmergic field is different, in that not only is the recipient undetermined, but so is the time when the action upon the field exerts its influence on a subsequent event. Time, and particularly sequence, is critical.

In contrast to networks defined by point-to-point links that represent the possibility of communication, or networks defined by communicative events, a stigmergic system is defined by a field on which events may have long-lasting effects that are observable at a later time. There is a conceptual network connecting the agent, the remnant effect of the event, and the (possibly many) observers of the event. The relation of a book author to the readers provides an everyday example. However, no conceptual map of links for possible communication makes much sense in a stigmergic field. The stigmergic network is defined by the observation of the influence of prior events, not by the traffic over defined node-to-node links.

Whether quasi-permanent, event-based ephemeral, or stigmergic, networks usually are interesting only in connection with events. Event sequence and timing are important in the dynamic behaviour of the network, and therefore in any visual representation of the network.

The mere existence of all these possible characteristics of networks illustrates that the display of relationships offers substantial challenges. This challenge has been passed to a new Technical Team, IST-059, which has the mandate to study the visualisation of networks.

2.2 USING THE VISTG REFERENCE MODEL IN PRACTICAL SITUATIONS

The explicit intention of IST-021/RTG-007 was that the VisTG Reference Model would be used by one or more nations in evaluating several systems proposed by other nations. For a variety of different reasons, most of the proposed systems were in fact not made available for evaluation. Three Slovak systems were proposed, but one was found to be unreleasable for security reasons while the other two could not be used without translation of the manuals. A UK system was expected to be releasable, but in the event, the release procedure took longer than the life of IST-021. Overall, five Nations proposed seven systems for evaluation, but no systems were evaluated outside their nation of origin. Only the Canadian VITA system was examined in any detail using the VisTG Reference Model. The Norwegian delegate was given a three day introductory course on the Norwegian tri-service NORCCIS II system and free ad hoc hands-on evaluation thereafter, but the available time did not match the complexity of the problem.

The comments of the Norwegian Delegate on using the VisTG Reference Model with NORCCIS were:

> NORCCIS II was a natural choice for a Norwegian system to investigate since it is the national system used on the operational level containing advanced visualisation for a number of tasks. I had a general interest of becoming more familiar to the system, and those responsible for its fielding and further development were quite interested in suggestions from IST-021/RTG-007. They offered a stand alone version on a laptop to the Paris meeting in May 2001, and they offered the group availability to their education center near Oslo in Norway. I attended a three days introductory course on the system in August 2001. The "exam" was to do 25 related tasks that had been part of the course; I completed only half of them in the one hour time available. The experience from this course supported several of the worries/problems formulated in Appendix 1 and was a background for some of my suggestions in the procedure for "evaluations" in Appendix 2. [Attached as Appendix 1 and 2 to Annex A of this report].
>
> I had some problems using NORCCIS II due to my rather poor knowledge of the military activities and concepts used in the detailed tasks. However, my main problem using the system was recalling the right interactions hidden in the menus/buttons hierarchies. This seemed as a common problem well known to those responsible for the system. They were aware of improvements, i.e. through more "direct manipulation". However, this was unrealistic due to the cost of redesign in the light of the many higher priority improvements continuously being implemented. I was surprised of their awareness of possible visual improvements; they seemed well informed about users' problems and suggestions. They did not really seem to lack good ideas, rather resources of implementing them!
>
> I was not able to "evaluate" NORCCIS II according to the VisTG Reference Model and the related procedures we have. I really lost the battle when thinking about how to do it in a meaningful way. "[The] 6 questions" immediately produced a lot of answers, but I could not organize the lot of them in the structure of the three level model. When planning an "evaluation" according to my own suggestion, the lot of organizing and work in relation to the perceived benefits made me lose that battle too.

Overall, these comments are not very encouraging for the use of the VisTG Reference Model in a quick appreciation of a complex system by an untrained passive observer. Interaction with experienced and novice users was a part of the evaluation, but the method discovered only issues that were already known to the users. The comment that may be most central to consideration of the method is probably in the final sentence of the quote, that the mass of organizing overwhelmed the benefits. Perhaps, in hindsight, this should not have been unexpected, since the method inherently produces a proliferation of questions, answers, and relationships, all of which must be combined in generating the critical evaluation of the system. Automation is an obvious suggestion to improve that issue.

2.2.1 An Evaluation of VITA

The following sections are largely based on a report prepared for IST-021/RTG-007 on the evaluation of two installations of VITA Delta. VITA is an evolving system, and the Delta version was its fourth incarnation.

An epsilon version exists, but was not available for this exercise. Henceforth, VITA Delta will most often be called simply "VITA". The main objective of the report was not to determine the effectiveness of VITA as such, since improved versions exist. It is to assess the use of the evaluation procedure using the VisTG Reference Model.

The first installation, on which this appreciation is mainly based, was on a Virtual PC (Version 5 by Connectix) running Windows 98 on a Dual processor 500 Mhz Macintosh over Mac OS 9.0.4, 9.2.2 and OS X 10.2.x. The second installation was on an Intergraph computer running Windows NT. This second installation was viewed only briefly, to determine whether differences other than speed were apparent. No such differences were observed, except that with the installation on the NT system the wheel on the two-button wheel mouse appeared to be non-functional because the Operating System lacked the necessary driver. The wheel did, however, function properly on the mouse used for the Virtual PC installation.

2.2.1.1 The VisTG Reference Model

The "VisTG Reference Model", is described under the name "IST-005 Reference Model" in the final report of NATO IST-013/RTG-002. Its use in evaluation is described above. It implies that there should be a series of steps in the conduct of an evaluation These are:

- Describe what the user wants to achieve, with a view to assessing his or her ability to perceive the degree to which it has been achieved and his or her ability to act towards achieving it. This goal can be at any level of abstraction desired by the analyst.
- Examine the degree to which the user can perceive what is necessary.
- Examine the degree to which the user can choose the necessary actions.
- Examine any impediments or external influences that affect the user's ability to perceive what is necessary.
- Examine any impediments or external influences the affect the user's ability to act effectively.
- Examine the degree to which other parallel events or tasks alert the user that it might be valuable for the user to attend to them instead of this task.
- Apply steps 1 to 6 to any parallel loops and recursively to lower-level loops whose existence is implied by the results of the earlier analyses. Higher level goals may also be implied; they are the reasons why the initial goal exists. If the higher-level goals are within the realm of the study, recursively use the six questions to investigate them.

It is not possible to perform all these steps properly with VITA on the basis of a brief examination, which is why the report title originally contained the word "appreciation" rather than "evaluation". There are two major reasons a full evaluation could not be done with the resources at hand. One is that VITA Delta is an experimental system that allows a large number of parameter variations in the manner of display, the effects of which may not be independent, leading to an astronomical number of possible combinations that might need to be tested. The other is that VITA functions most effectively when run on the fastest possible machine. Although Virtual PC seems to emulate the functions of a real Pentium very accurately, it is an order of magnitude slower than the real PCs now available. Furthermore, the Virtual PC does not take advantage of specialized 3D graphics hardware, for which the program is optimized.

To some extent, the relative slowness of the Virtual PC helps the appreciation, because one of the issues is a problem of scaling that appears with an increase in the number of displayed entities. On a real PC with hardware acceleration for 3D displays, this issue does not appear until a large number of entities are shown, but with the Virtual PC, which does not permit hardware 3D acceleration, the problem shows up with a few tens or low hundreds of displayed entities. The scaling problem would, however, be an issue in many potential uses for VITA.

2.2.1.2 VITA

VITA can be considered from more than one viewpoint. At heart, it is a concept for displaying interactively the many-to-many relationships between, on the one hand, questions that might be posed implicitly or explicitly by a user and, on the other, a universe of data within which the answers to those questions might be found. As implemented, VITA Delta exists within a complete application, the purpose of which is to allow the user to discover relationships that determine the structure of a Web site, or to help a user refine a search for the answers to specific questions whose answers might be discoverable from data within the Web site.

The implementation of VITA within a complete application necessarily requires decisions as to the details of the display. Different incarnations of VITA have allowed the user different freedoms of choice among display possibilities. For example, in one version, the user can observe the relative locations within individual Web Pages at which elements responsive to the queries reside. VITA Delta does not have this display option, but it does allow many choices of colour, shape, and transparency of the displayed entities and the ways those entities can interact with the user. Most of this appreciation uses the default values of those choices.

VITA can be regarded from a narrow or a wider viewpoint. On the narrow view, it is a way of discovering within a single Web site those pages that do or do not contain specific subsets of a set of keywords provided. This is the capability provided by the installed version, which relies on the Google search Engine as one of its sources for its data. Google only returns the existence or non-existence of the sought keywords within the page; this limitation does not allow the local user to perform detailed structural analyses of the relations among those keywords within the Web pages. A second search Engine known as “maglia” was available for use within one site, but this was little used in the present evaluation.

In the wider view, VITA is a concept for displaying a universe of objects that have attributes, so that the objects with similar relationships among the attribute values are clustered together in a 3D space. The VITA Delta implementation uses a Web site as the universe, its pages as the objects, and the Boolean value of the existence or non-existence of selected terms as the measure of the affinity between the objects. Some screenshots of VITA are shown in Figure 2-4.

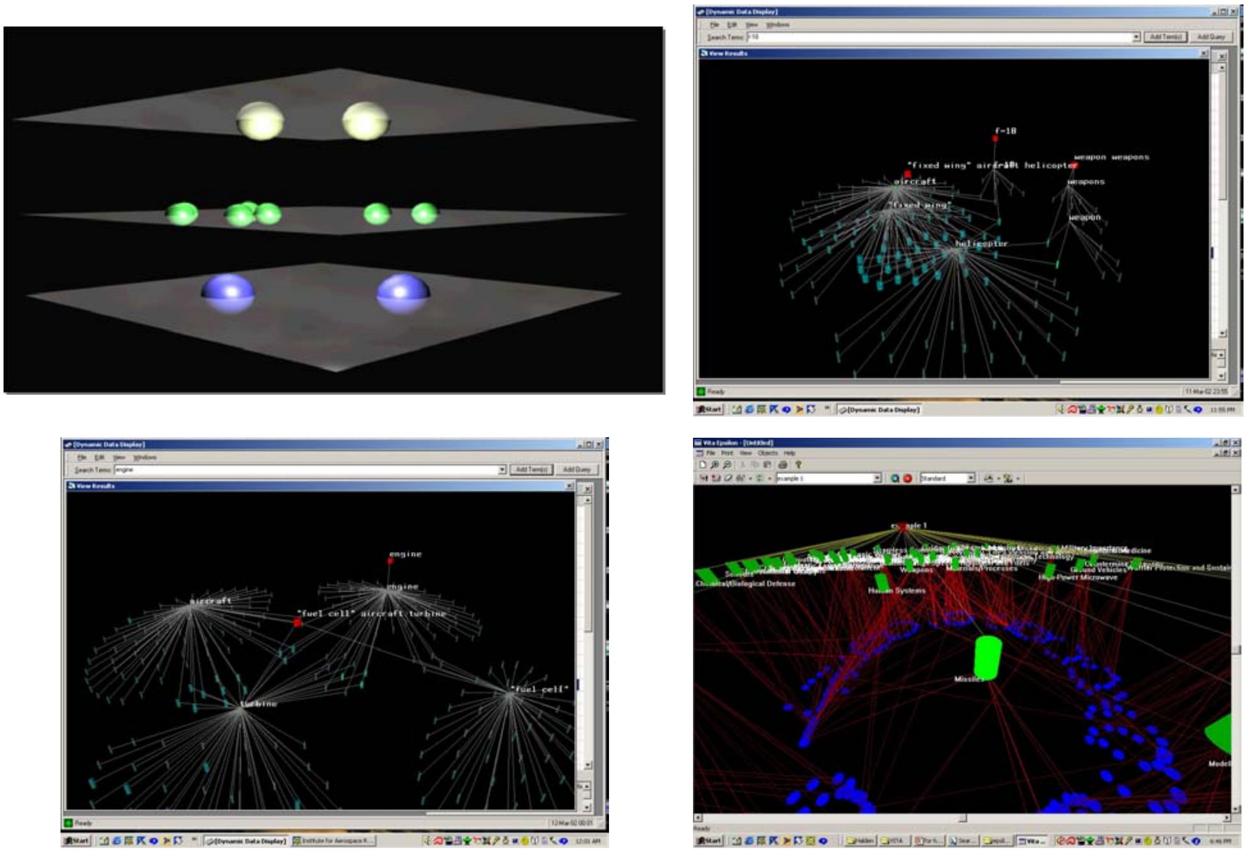

Figure 2-4: (Top left) The Canonical Three-Layer Structure of VITA, with "Queries" in the Top Layer, "Concepts" in the Middle, and "Documents" in the Bottom Layer. (Other panels) Various screenshots of different VITA displays.

In either the wider or the narrower view, VITA should be considered as an "Engine" within the VisTG Reference Model, not as the application within which the Engine is implemented. An Engine, in the VisTG Reference Model, is an intermediary between the dataspace of interest to the user and the display systems that the user sees (or feels, or hears, or speaks to...). VITA Delta exists within the context of an application that aids searching Web sites, but VITA itself is not a means of searching Web sites. Any evaluation must tease out those limitations that are imposed by VITA from those that are imposed by the application context in which it is used – for example, the Google limitation that disallows the discovery of the proximities of keywords locations within Web pages.

In its current implementation, VITA includes not only the Engine that computes the desired relationship patterns, but also the Input-Output functions that create the display and that allow the user to influence the view of the display. These I/O functions are affected by choices that can be made interactively by the user. The VisTG Reference Model evaluation, therefore, should consider at least two loops: the Engine loop and the I/O loop. However, these can be separated in considering the usefulness of the VITA Engine for more general applications.

2.2.1.3 VITA Concept

VITA is conceived within a universe of text documents, of which Web pages are a convenient example. It could, however, be extended to other kinds of entity, either textual or not. During the September 2002 IST-036/RWS-005 Workshop, for example, one of the syndicates raised the question of how to allow a commander or the staff to visualise the potential sources of relevant time-critical data in coalition operations, when the different allied forces have different restrictions on what is releasable to whom and with what delay. VITA seems potentially valuable for helping the command staff with this task. The different national or individual data sources would take the place of the Web pages with which VITA Delta deals, and the queries would be the commander's description of the kinds of information that would help his/her decision-making. Such queries are sometimes given the label "directed telescope", or rather, they constitute the directions to the telescope.

For VITA, considered as an Engine rather than in its current instantiation, the primary requirement on the dataspace is that there be a computable commonality among the discrete data packages. In VITA Delta, the commonality is the pattern of occurrence or non-occurrence of selected terms in the individual Web pages. In what follows, the "discrete data packages" will usually be documents or Web pages, but they might just as well be the constellations of information the command staff has about the information types that might be available from, and the information-handling policies of, allied units and other potential information sources. Nevertheless, they will be called "documents" or "pages" in what follows.

Given a measure of commonality or similarity with respect to the selected attributes of the documents in the dataspace, the VITA Engine creates a self-organized structure of the document representations so that documents with high similarity are represented as neighbours in a two-dimensional space. The "gravitational Engine" that allows the documents to attract and repel each other according to their degree of similarity in the Boolean pattern is the core process of the VITA system. In some versions of VITA, the representations of the documents can be probed to determine where in the document the key terms occur, and in VITA Delta, the documents are Web pages that the user can request be opened in a conventional browser.

2.2.1.4 Using the VisTG Reference Model for a VITA Appreciation

VITA can be considered as a program that is installed on a computer, or it can be considered as a concept that has been implemented in one of various possible ways. As a concept, it involves two levels of interest: defining relationships among entities in a dataspace, and displaying those relationships. As an implemented program, one must consider details of the interface which may not be essential to the VITA concept. Analysis using the VisTG Reference Model can be treated from either viewpoint.

Evaluation using the VisTG Reference Model can be initiated at any point in the system being evaluated, by determining a purpose that is being served at that place in the system. That purpose defines a loop that includes what the user needs to be able to perceive in order to achieve the purpose, and what actions the user may perform to achieve it. Having defined one purpose, the evaluator may proceed "outward" by asking why that purpose is necessary – in other words, what higher-level purpose it supports – or "inward" by examining what subsidiary purposes serve to enable the perceptions and actions that are needed for the outer (higher level) purpose. The answers to the questions, as with any appreciation using the VisTG Reference Model, imply questions to be asked both of the inner loops or lower levels of abstraction, and of the more outward loops that provide ever more complete answers to the question of "why is this wanted?".

For the present appreciation of VITA, the initial purpose is taken to derive from the user's task, which requires the user to discover and to examine entities that have certain properties in a discrete dataspace. In the implemented version available for examination, the dataspace consists of a Web site, the entities are Web pages, and the properties or attributes are the occurrence or non-occurrence of specific words (search terms) in each Web page on the site. From time to time in this appreciation, reference is made to a hypothetical command-post application, in which the command staff in a coalition operation needs to determine which governmental and non-governmental sources of information are likely to provide the quickest, most relevant, and most reliable information concerning an evolving situation.

2.2.1.4.1 Level 1. Outer Loop Questions

Q1. Define an initial loop by asserting what the user wants to achieve.

Instantiating the question for the particular loop to be evaluated: Defining any loop depends on identifying what the user wants to achieve at that level of analysis. The questions for the rest of the loop concern to what degree the information and actions available to the user allow this to be done.

Answering the question: For VITA, the outer loop is concerned mainly with the concept. For this appreciation, we assume that the purpose is for the user to discover and examine more closely entities with desired but non-obvious characteristics (Web pages relevant to particular topics, in the implemented version), and that relevant entities can be approached by means of navigating through patterns of attribute similarities among the entities in the dataspace.

An alternate answer is that the purpose is to help a commander in a coalition operation to determine the most likely sources for acquiring accurate, timely, and relevant information in an unexpected situation. The underlying assumption is that the Engine has access to a database of information sources, which includes such data for each source as the accessibility of different kinds of information (not all nations release all information to all other nations), the past reliability of information from that source, the time it usually takes to access information from that source, and so on.

Comment: Instantiating and answering this question is probably the most crucial aspect of an evaluation (or a design) using the VisTG Reference Model. It may be difficult initially to avoid mixing levels in defining what should be a single-level loop, and particularly at the higher levels the purpose of the loop may be quite ill-defined. That it is ill defined may be intrinsic, but the more precise the definition of the purpose, the easier it is to answer the subordinate questions that deal with the means of implementing the purpose and the possible impediments to success.

Q2. Can the user to perceive what is needed?

Instantiating the question: The intended result of using VITA is that the user can see the different patterns of relationships among the entities in reference to the selected attributes, and then to be able to select for closer examination entities with desired relationship patterns from the dataspace. Therefore the user needs to be able to perceive the groups of entities that are similar to one another, and to be able to see enough detail to extract individual entities that seem worth investigating further. further, the user needs to be able to see which attributes appear to be discriminatory, and to be able see what is needed in order to add and remove attributes so as to approach the most relevant entities in the dataspace. In the current implementation, the entities are pages within a Web site, and the relationships are determined by the occurrence and non-occurrence patterns of search terms in the pages.

Answering the question: Can the user perceive what is needed? The main VITA display consists of an arrangement of three parallel planes in a 3D Euclidean space, with nodes representing queries (or requirements for types of information) in one plane, nodes representing concepts (search terms, or components of the required types of information) in a second plane, and nodes representing entities (Web pages, or coalition and non-governmental information sources) in the third plane.

The nodes of the middle plane (representing concepts, or information components) are connected by lines to the top-plane nodes for queries that included those concepts, and to the bottom-plane entities (Web pages or information requirements) that contain them. Pseudo-gravitational attraction is used to locate nodes within their specific planes, and the dynamics of the pseudo-gravity arranges the nodes so that like clusters with like. In principle, this should allow the user to perceive what is necessary in order to select and examine entities with particular patterns of attribute values. In practice, other considerations may make it more difficult for the user to achieve the desired perceptions. That aspect of the evaluation belongs at the next inner level of analysis. At this level, it suffices to say that the information is, in principle, available. The implementation, however, does not make it obvious without instruction that the information required for adding and subtracting attributes in the queries is available.

In the current implementation, which is concerned only with Web pages, the Google search Engine is an option for searching a Web site for relevant pages. Other options are also available, but were not examined in this appreciation. The Google Engine returns only whether a term existed in the Web page, not its location. This means that the proximity between occurrences of the search terms in the delivered pages is not available to the VITA Engine. To have such proximity information made perceptible would greatly enhance the user's ability to fulfil the purpose of discovering highly relevant pages.

Comment: It is tricky to instantiate and answer this question only at the level of the particular loop under consideration. The temptation is always to incorporate factors from lower-level loops, and it is not always obvious when this is happening.

Q3. Can the user act to influence what is perceived in relation to the purpose of the loop?

Instantiating the question: The user has to be able to specify the attribute values that define the relationships that will select the entities to display. Once the entities have been displayed so that the relationship patterns are evident, the user has to be able to navigate within the display by adding and subtracting concept search terms, and by changing the viewpoint onto the 3D display, so as to allow examination and extraction of potentially interesting entities. In the current implementation, the user needs to be able to navigate so as to select individual pages for display in a browser. In the hypothetical information-source discovery application the user needs to be able to refine the description of the kinds of coordinated information types that are required, in response to initial discoveries from the initial displays.

Answering the question: At this level, the evaluator asks whether the user has the manipulative facilities to provide the Engine with the information necessary to generate the needed displays. In principle, those facilities do exist in VITA, but examination in more detail at the next (inner) level may show up deficiencies or difficulties. So, at the next level, the evaluator must ask whether the software allows the user to manipulate input devices that allow the displays to be controlled effectively.

Comment: Again, it is tricky to instantiate and answer this question within the single level. For VITA in particular, the tendency is to consider the ability to control 3D rotations and to select entities in a cluttered 3D space. Such questions belong at the lower level. At this level the issue is whether the user has the information or control that would be required to perform the rotations or selections, and the answer is "Yes."

Q4. What impediments might affect the user's ability to perceive what is necessary?

Instantiating the question: The question intrinsically depends on the answer to the question about what the user needs to be able to perceive. The user needs to be able to see a representation of each individual Web page, or to be able to navigate so that this can be done. Are there impediments to being able to do so? Furthermore, the user needs to see whether representations of Web pages relevant to the purpose exist in the display, which means that the representation of each Web page needs to have some indication of why it is being displayed, relative to the topic of interest to the user. One potential impediment might be the inadequacy of the display of individual Web page representations. Another might be the possibility that the display of one entity might impede the visibility of another.

Answering the question: If the Engine selects too many entities from the dataspace, there might be a problem of clutter in the display. Clutter presents two potential problems. One is that the user may not see the truly relevant entities among the mass of less relevant ones. The other is that the display of the relevant entity may be masked by the display of one or more other entities. In the narrow view of VITA this means that it is important that the Web site being examined should not be too complex, or that not too many terms be defined for selection, since every page that contains any of the keywords is made available for presentation. In the wider view, this problem might be exacerbated. If the entities are all shown as alike, the user may have a difficult problem to discover which ones are really interesting, as opposed to those that adventitiously have a particular interesting pattern of attributes.

Comment: As before, the temptation is to mix levels. By looking at the details of the display, items that are intrinsic to the issues of this level may be overlooked. In VITA, the issue at this level is not so much why particular problems may arise, so much as what elements of the task would be likely to give rise to problems.

Q5. What impediments might affect the user's ability to influence the Engine appropriately?

Instantiating the question: The actions the user needs to do are to navigate within the display, to select individual entities (Web pages or concept nodes), to determine whether the entity representations are worth deeper examination, and to pass them to other software for that further examination. From time to time, a concept node may need to be removed, or a new one added, as part of the process of navigating through the dataspace.

Answering the question: Impediments to these actions might include insufficient degrees of navigational freedom, inability or difficulty in selecting the desired entity, difficulty in rapidly determining whether the selected entity is worth deeper examination, and difficulty in performing that deeper examination (using other software – in the case of the implemented version, that other software is required to be Internet Explorer 5.5).

Comment: If impediments exist, there are no obvious implications for design improvement at this level of analysis, though the potential impediments do set up questions that must be addressed at the next lower level.

Q6. Is there any provision for alerting the user to parts of the dataspace that should be expected to be of interest?

Instantiating the question: One might consider that the main objective of VITA is to provide a display that aids the use to find entities of specific characteristics, and therefore of potential interest--in other words to be an alerting system. However, that is not the sense of Question 6. Question 6 addresses the automated analysis of parts of the dataspace so that those that have characteristics of potential interest can be displayed (perhaps auditorily or haptically) in such a way as to draw the user's attention. The concept of "alerting" implies differential display of "interesting" parts of the dataspace.

Answering the question: In this sense, VITA does not employ alerting.

Implications for design improvement: There is no obvious way for the user to inform the system as to what would make a Web page "interesting", and failing that knowledge, it is not possible for the system to provide an alerting service.

2.2.1.4.2 Level 2. Inner loop questions

Q2.1. Define a next-level loop

Instantiating the question: For any visualisation system, certain support Engines are likely to be required. Characteristically, among these are:

1) Navigation Engine: Navigate to the relevant part of the dataspace;
2) Data Selection Engine: Select the relevant data;
3) Algorithm Selection Engine: Choose algorithms to apply to the selected data; and
4) Various Engines: Execute the appropriate algorithms on the selected data.

For VITA, only the Navigation Engine and the Data Selection Engine seem relevant. There is no algorithm selection, at least insofar as they refer to algorithms to be applied to the data, since the core concept of VITA involves a single algorithm. However, VITA does allow the user to make many choices as to how the data are displayed. In a sense, this choice implies a procedure for algorithm selection, but it might be better to consider it a separate class of Engine, which might be called a "System Control" Engine. This aspect of VITA will not be discussed in this evaluation, as for most purposes default values of the control parameters were used. The effects of varying those parameters would require extensive experimentation, which is outside the scope of this evaluation.

At this level, Data Selection means the selection of individual entities (Web pages) from among those initially selected at the outer level as being of interest and therefore made available for display.

Navigation is important, especially when there are many displayed entities, but Data Selection is trivial, being (in this implementation) only a selection by point-and-click. Because the Data Selection is so trivial, the Navigation must be carefully done. Navigation is the only means whereby data can be displayed in a way that allows for point-and-click selection, whereas the provision of more powerful selection tools would take some of the load off the Navigation Engine.

Answering the question: The purpose of any inner loop is determined by the answers to Questions 2 or 3 of the loop immediately outside it. In this case, we have two answers.

a) The purpose of one inner loop is to create one or more displays that allow the relationships among entities, as defined by the specified attributes, to be perceived. The outer level loop has generated the data according to the user's purpose at that loop level. At this inner level the generated values for the relationships must be formed into visual patterns on the display. Navigation occurs among those patterns. Navigation is unimportant when there are only a few tens of entities (Web pages), representations which can be placed on a screen without mutual interference and without excessive visual overlap in at least some views. However, when there are many entities to be displayed the ability of the user to see where interesting entities might exist and to be able to focus on them becomes harder to achieve. Good navigation facilities become very important. So, the purpose of this loop is to allow the user to navigate to the most relevant entities.

b) A second inner loop allows the user to specify the attributes that determine the relationships to be plotted. In other words, it allows the user to specify the queries in the Web-page implementation, or the specifications of the potentially useful information sources in the hypothetical command application. In the narrow view of VITA, as implemented, this is a trivial keyboard entry of terms that must be included in any Web pages to be selected, and so this loop will not be analyzed further. However, in a wider view of VITA the way the user defines attributes might be an issue that requires evaluation.

Q2.2. What does the user need to perceive?

From this point on, I assume that the approach to instantiating the question is, if not self-evident, at least implicit in the answer to the question.

Only loop (a) above is under consideration.

In either the wide or the narrow view of VITA, the relationship patterns are reflected in the spatial separation of the entity representations on a single plane in the 3D display space. A second plane contains the individual attributes as nodes located symmetrically around one or more rings, but the locations of the nodes on the ring signifies nothing about the relationships among the attributes themselves. A third (top) plane consists of nodes that represent individual "queries" or in the hypothetical command application, the specifications of the kinds of information that would be useful for the task of the moment.

VITA Delta allows the user to restrict the display of entities to those that contain at least N items from any query. This ability allows the user to constrain the displayed entity nodes (representing Web pages) to those that are likely to be at least in the correct contextual domain. For example, if the query includes the word "bank", the pages reported by Google would include those about financial institutions as well as those relating to the ecology of riverine systems – but if the user elected to display only those pages that contained at least three of "bank" "ecology" "river" "lake", it is unlikely that any of the Google-reported pages containing financial institutions would be displayed.

In the narrow view, queries are selection instructions to the supporting search Engine (limited to Google for this appreciation). Different queries may use overlapping search terms, and even when search terms do not overlap, different queries may select overlapping sets of documents. If either case occurs, the overlap is signified by the display of connecting lines between the cylinder representing a document and the balls representing the concepts that occur in that document, and between the concept balls and the blocks that represent the queries that contain those concept terms. (In the present implementation, a "concept" is equivalent to a word, rather than to its meaning).

The relationships among the documents selected by different related queries, added to the user's election to display only pages containing N of the search terms, further enhances (in principle) the ability of the user to focus on the truly relevant Web pages.

Navigationally, the user needs to be able to shift focus within the display of entities (Web pages), attributes (concept or query terms), and queries, so as to examine in more detail the individual entities (drilling down).

In general, there will be more entities than can be shown in a window of realistic size while maintaining a useful separation among the different entities. This implies that one of three approaches to creating the display must be used: either

1) The user must always be able to see enough of the attribute pattern information to allow scrolling the display to the region containing the interesting entities,
2) The display must be of the "fisheye" type, in which the items of interest are shown in greater detail than the others, and the user can move the focus readily to the region of immediate interest, or
3) There must be a means of coalescing a group of similar entities into a larger-scale display unit.

Typically, in the "fisheye" context, greater detail at the focus is accomplished by expanding the central part of the window area with respect to the periphery, with the degree of expansion being sufficient to allow the entire set of extracted data to be displayed at the same time. This technique could be used with VITA, with a moderate number of entities, but in the present implementation, the space is Euclidean, meaning that entities often go offscreen, removing them as parts of the context, as well as sometimes making it difficult to see which entities are linked to which concept.

Q2.3 What actions does the user need to perform to allow the necessary things to be perceived?

The user must be able to orient the display so as to see the entities that have interesting sets of attribute values, and to shift the viewpoint so that the interesting entities can be individually distinguished and selected. Having selected an entity, the user must be able to determine quickly whether it is worth devoting more attention to that entity, and to be able to request greater detail if the quick look makes that seem desirable.

Provision is made in the implemented version of VITA for a rudimentary form of this quick-look capability, mousing over a document representation ("brushing" it) causes a few words of the Web page to be displayed in a text field above the document. (Note: there is a bug in the implementation, in that even when the cylinder representing the document is at the top of the display space, the small text sample appears above it, meaning that it is off-screen and invisible). However, the brushed text is often not enough to allow the user to determine whether it would be worthwhile to use the much slower browser method of seeing the entire Web page. However, no such "full disclosure" display seems possible in the hypothetical command application.

Q2.4 What impediments might prevent the user from perceiving what is necessary?

The display might make it difficult to see which entities have interesting attribute patterns, or to see to select among those that do. It might not be easy to see the initial more detailed view of a selected entity. It might not be easy to see where to navigate within a larger dataspace.

In the current implementation, all documents are represented by similar cylinders that have no indication of their content except that their diameter varies with the number of search terms that exist somewhere in the document. The only indication of relevance is the existence of visible lines that connect the Web page representation to the concept nodes that represent the terms contained in the document. If those concept nodes are off-screen, which can happen when the display is zoomed to reduce clutter, it becomes hard for the user to identify which concepts are present in any particular Web page.

Implications for design improvement: In the Web search version of VITA, this means that the pages should be displayed with some differentiation based on content other than simply the fact that they have a particular Boolean pattern of containing the search terms used in the queries. This is not done in VITA Delta, though examples have previously been shown from earlier versions of VITA in which documents were shown as cylinders that vary in size or in internal structure. That display or something similar could ease the task of discovering truly relevant Web pages. If (as is the case with the Google search Engine) the information simply is not available, then no design improvement can help.

However, something could be done to improve the possible situation in which lines cannot be identified with the offscreen concept nodes to which they are connected. Without specifying the mechanism (which is a problem at a lower level of analysis), a design improvement would be one in which the concept nodes could be identified whatever the level of zoom the user chooses to reduce clutter.

In the implemented version of VITA, problems arise when there are many displayed documents. When that is the case, not only do the entities cluster too closely to allow easy selection, or even easy visibility, but also the displayed lines that link the three levels (queries, terms, and documents) often obscure the documents displayed. Even if the display is rotated so that the view is from beneath the document level, the lines may obscure the concepts. The lines are required so that the user can trace from a concept to an associated document or from a document to the various terms that it contains, but at the same time they are a hazard when the user is trying to select a document.

Associated with the above comment is a mild criticism of the imposed limitation on the size of the display space allotted to the entire VITA system in this implementation. The largest window allowed is only 1024 x 768, in which the3D display window must be shown along with other required windows, even when VITA is running on a 1600 x 1200 screen that has nearly 2.5 times that area. When the displayed documents are both tightly clustered and extend beyond the boundary of the 3D display window, this artificial limitation is difficult to accept. It is possible to expand the 3D display window size by dragging on the window corner in the standard manner for such a window-based operating system interface, but when one does so, the expanded lower portion of the 3D display remains blank.

Both techniques (1 and 2 above) would be difficult to implement effectively with the kind of display currently used by VITA if the number of displayed entities were to grow to the thousands or millions, but a flat (Euclidean) display mainly contained on a single screen works reasonably well when there are less than the low hundreds of entities. However, the lack of a means to display distant context while a particular focal region is under close examination might prove an impediment to the wider applicability of the VITA approach. On the other hand, some arrangement for automatic or manual hierarchic grouping of entities might permit the VITA approach to be retained even with very large numbers of entities.

In the current implementation, all documents are represented by similar cylinders that have no indication of their content except that their diameter varies with the number of search terms that exist somewhere in the document. The only indication of relevance is the existence of visible lines that connect the Web page representation to the concept nodes that represent the terms contained in the document. If those concept nodes are off-screen, which can happen when the display is zoomed to reduce clutter, it becomes hard for the user to identify which concepts are present in any particular Web page.

Implications for design improvement: In the Web search version of VITA, the pages should be displayed with some differentiation based on content other than simply the fact that they have a particular Boolean pattern of containing the search terms used in the queries. This is not done in VITA Delta, though examples have previously been shown from earlier versions of VITA in which documents were shown as cylinders that vary in size or in internal structure. That display or something similar could ease the task of discovering truly relevant Web pages. If (as is the case with the Google search Engine) the information simply is not available, then no design improvement can help.

However, something could be done to improve the possible situation in which lines cannot be identified with the offscreen concept nodes to which they are connected. Without specifying the mechanism (which is a problem at a lower level of analysis), a design improvement would be one in which the concept nodes could be identified whatever the level of zoom the user chooses to reduce clutter.

Above, it is noted that selection of truly relevant Web pages would be eased by the fact that different queries do or do not point to the same Web page. However, the display does not make this relationship apparent. If two queries contain the same concepts, the display shows only the link between the page and the concept. The usefulness of the VITA Engine might be substantially enhanced if the display were to show even for how many queries the page is a response. This might be done by a variety of means, such as showing the displayed cylinder as a set of disks, one for each responsible query, or by allowing for the optional display of lines linking pages directly to queries rather than only to concepts. Such extra lines would greatly clutter the display, if shown all the time, so their display would have to be controlled by the user in real time (e.g. by a specific mouse action, by an on-screen button, or by specific keystrokes).

Q2.5 What impediments might prevent the user from performing the required actions?

Insufficient degrees of freedom might be available for shifting the viewpoint within the 3D dataspace, for selecting an entity out of the clutter of displayed entities, or for requesting the display of preliminary detail about the entity. The implementation of the input interface devices might be insufficiently precise, or there might be excess delay between the user's actions and the effects of those actions on the display.

The three major impediments to action are ordinarily the three mentioned above: limitation of degrees of freedom, imprecision, and timing inaccuracy. In addition, unfamiliar mappings between user actions and the effects on the display also are a frequent cause of problems. Unfamiliar mappings can often be learned, if they are not too strange, but the other problems are intrinsic, and must be fixed within the computer, not within the human.

In the implemented version of VITA, control is through a two-button wheel-mouse. The interface allows the user to control the bindings among mouse actions and changes of viewpoint on the 3D display. The default key bindings allow the necessary actions, but I did not find them very intuitive. That problem is one that is easily overcome by learning. The fact that the mouse allows only a few degrees of freedom is not.

Implications for design improvement: It is conventional in window-based displays to provide scroll-bars along the X and Y dimensions of the window to allow the X and Y pan, one example of changing the viewpoint onto the 3D display.. If this were done in VITA, even a one-button mouse would be able to achieve a 3D pan without compromising other possibly useful degrees of freedom for control of the display, such as brief examination of different aspects of the contents of an entity, or making and expanding groups of entities. Accordingly, I strongly recommend that any future development of VITA incorporate conventional scroll-bars in the window in which the 3D display is shown. Alternatively, manipulation of the 3D display viewpoint might be achieved by controls in the window-frame in the manner of many VRML viewers.

Q2.6 What provision is there for alerting the user to "interesting" parts of the dataspace?

In VITA as implemented, there is no provision for alerting, and it is not clear how such a provision could be made. Alerting requires that criteria for alerting should be available to the alerting Engines. Those criteria might be hard-wired (as is the colour change of a recently selected entity in the implemented VITA Delta), or they might be communicated by way of the textual or graphical interface.

In the Web-page version of VITA, one might imagine adding alerting criteria related, for example, not only to the existence of search terms in the displayed document, but also to the proximity of those search terms to each other within the text, if that information is returned by the Google search Engine.

2.2.1.5 Conclusions of the VITA Appreciation

This brief appreciation of VITA was based mainly on an installation on a Virtual PC running on a Macintosh computer. Installation was relatively easy, but needed real-time guidance. If VITA were to be made available to a wider audience, it ought to be fairly easy to create an installation procedure that did not need this guidance.

VITA appears to be worth developing further. One potentially valuable development would be to reduce its fairly tight connection with examination of the Web, and enhancing the provision of Application Programming Interfaces (APIs) to facilitate the use of the VITA Engine in a variety of applications. In other words, VITA should be recognized to be an Engine, rather than being treated as an application in and of itself.

The applications for which VITA seems potentially valuable are those in which user-created constructs (queries, in the Web application for which VITA Delta is designed) are related many-to-may with dataspace entities (Web pages in VITA Delta) through an intermediate form (concepts embedded in the queries and in the Web pages, in VITA Delta).

In an abstract sense, VITA is an instance of a translation device that works through a common intermediary language. “Words” in language A (queries, in the Web application) have a range of connotations that partially overlap “Words” in language B (pages in the Web application). Effective translation from A to B requires that the appropriate connotations be conveyed. For example, are the connotations of “bank” in context associated with water or with finance? If an intermediate language is available in which different connotative topic areas can be expressed, then translation done with reference to that intermediate language is likely to be more accurate than a direct translation.

In the Web application of VITA Delta, the manifest relations among the search words in a query take the place of the context in a translation. A query in which “bank” co-occurs with “river” is unlikely to be satisfied by a Web page in which “bank” occurs without “river” but with “interest rate.” Using the Google search Engine limits the ability of the VITA search Engine to make subtle connections based on the probabilities of neighbouring occurrences of the search terms, as it reports only a Boolean true or false for the proposition “term exists”. Even with this limitation, the VITA display facilitates the discovery of pages that have a higher than average likelihood of satisfying the information requirements implied by the several queries that may enter into any single display. In part, this ability is enhanced by the user’s ability to limit the displayed entities to those that contain at least N of the terms used in each initial query.

The appreciation used the VisTG Reference Model to suggest stages in the evaluation, and to identify a few places where changes in the implementation might improve the usability and usefulness of the concept. The primary difficulty seems to be with scaling to situations in which many documents and concepts must be displayed, and with the degrees of freedom to allow rapid, intuitive navigation within the 3D display space. The display might also be made more efficient if the initial representation of a displayed entity were to provide some indication of its content, rather than requiring the user to “brush” each document individually in order to discover whether it would be likely to be worth investigating further.

2.3 CONCLUSIONS ON USING THE VISTG REFERENCE MODEL

Both attempts to use the VisTG Reference Model to evaluate real systems gave the evaluators problems in trying to separate out the different levels of abstraction necessary for effective analysis of single feedback loops. The real complexity of the system structures also made it difficult to keep track of the cross-influences. Nevertheless, both attempts did find areas of potential improvement in the systems evaluated, albeit that in the Norwegian experience, those areas had been observed previously by other methods, particularly the daily experience of the users. In a way, this is an encouraging sign, though it would have been more encouraging had the evaluation shown a possible improvement that had not been previously noted, but that was accepted once found.

Both attempts showed that the method requires automation if it is to be useful in any but toy circumstances. There is simply too much data and too many relationships for a pencil and paper approach to be successful. Software could be developed that would embody the relationships inherent in the model and that would help guide the evaluator or designer, ensuring that potential areas of difficulty were at least signalled as needing attention.

No software or conceptual model can substitute for experience, either the experience of a designer that supports an expert intuition as to what will work and what will not work, or the experience of a user who has enjoyed or been frustrated by an existing system. However, a good model supported by good software should enhance the ability of both expert and trainee designers and evaluators. IST-021/RTG-007 was not able to conduct sufficient experiments to determine whether the VisTG Reference Model might be a basis for developing effective support software for design and evaluation of visualisation systems, and the question remains open.

Chapter 3 – EVALUATION OF THE UK MASTER BATTLE PLANNER

Margaret Varga, Don Brealey and Chris Walker
QinetiQ, U.K.

3.1 INTRODUCTION

The principal aim of an air campaign planning organisation is to achieve the objectives defined by a higher authority through the creation of a plan which makes the most effective use of the available resources against the nominated targets within the time limit.

The UK Master Battle Planner (MBP) provides a graphical tool for visualising the battle scenario and developing the air campaign plan [1]. The MBP is a map-based approach to planning which allows visualisation of the entire battle scenario, see Figure 3-1. It provides a dynamic real-time visualisation of the battle scenario and thus provides a high degree of situation awareness.

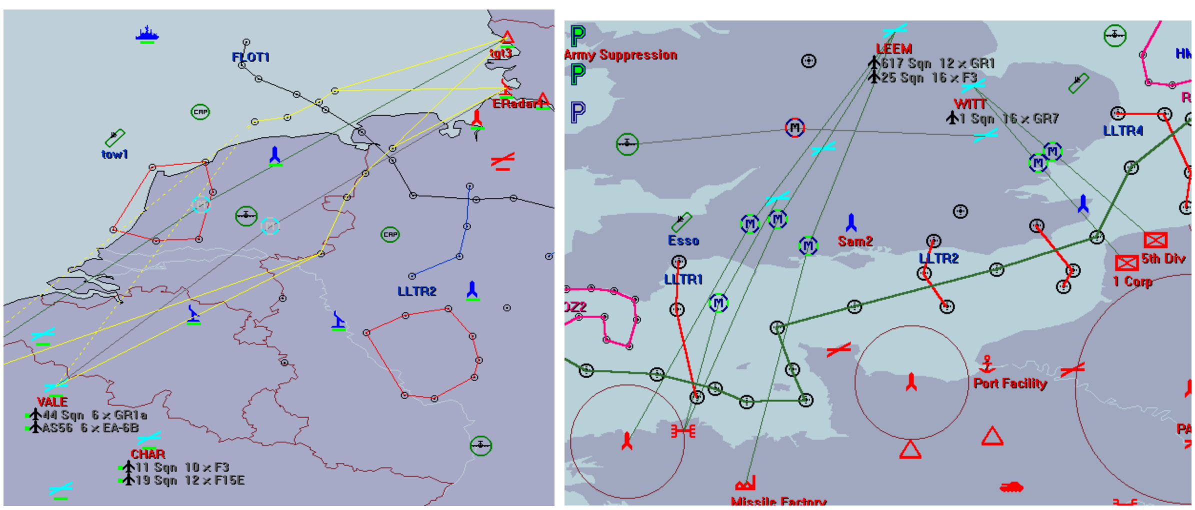

Figure 3-1: Direct Interaction with Map-Based Plan.

The MBP is designed to be flexible and intuitive, and can be used in a stand-alone mode or in conjunction with more complex planning systems. It supports interrogation and modification of all objects by interaction with them on the displayed map or in other data views including; Gantt charts, histograms, etc., see Figure 3-2.

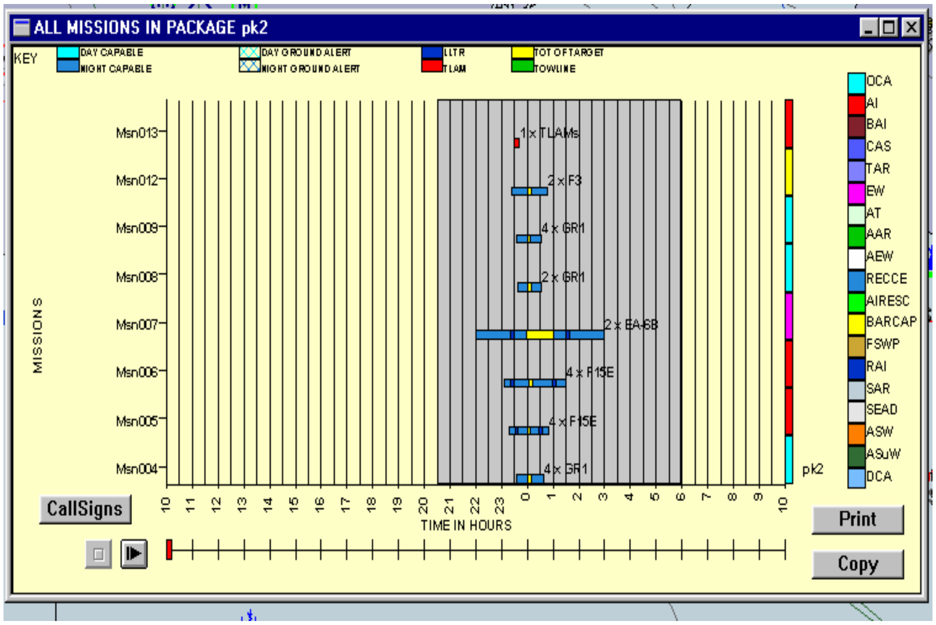

Figure 3-2: Direct Interaction with Mission Timings.

The MBP can be used to produce an initial draft Air Tasking Order (ATO) with packages and missions: it can then be used to analyze and refine the plan. The plan can be analyzed to show:

- Summary of missions in time.
- Summary of Air Unit tasking.
- Apportionment of resources by mission roles, with graph (Battlegram) of sorties by mission role over time.

The following is a list of MBP's capabilities:

1) **Scenario Visualisation** – provides the operator with facilities to establish and visualise the battle scenario with a toolset. There are several different types of visualisation view; map, Gantt charts, histograms, etc. The map view can present the entire battle scenario. The MBP is interactive and objects can be interrogated and modified in any of the views.

2) **Campaign Objectives** – provides the operator with the ability to refine the campaign objectives, i.e. based on the Commander's intent and guidance and to establish their connection to the nominated targets. A method of auditing the derived plan against the objectives is also provided.

3) **Mission/Package Planning** – supports the development of Offensive, Defensive and Support mission types with associated routes, etc. Using Windows-type displays, the planner can readily develop and modify complex Packages via interaction with objects displayed on the map and other data views, for example see Figure 3-3.

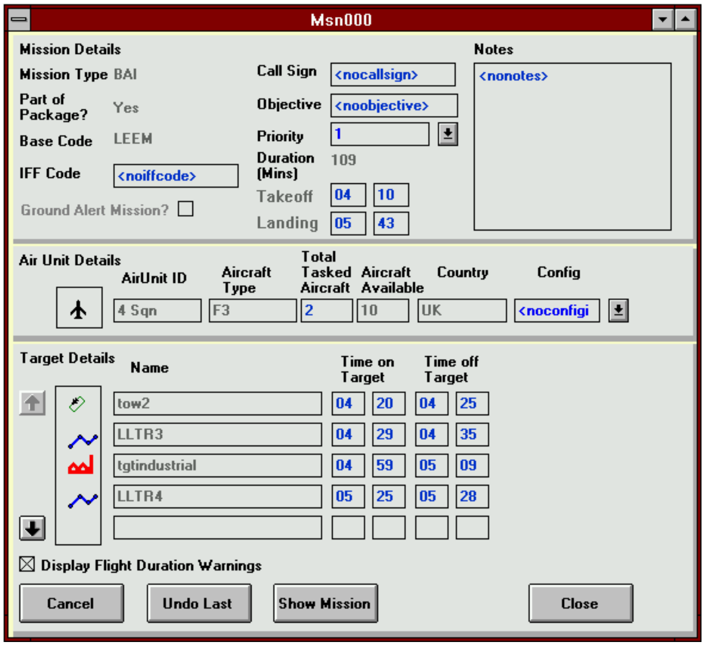

Figure 3-3: Mission Route Display.

4) **Plan Analysis/Refinement –** A suite of tools is provided to enable the planner to analyse the plan in terms of aircraft available and sorties tasked and to make changes where necessary. A 'fly-out' of the plan is provided to further support plan analysis.

5) **Briefing the Commander** – Significant time is spent preparing briefs for the Commander. With MBP, as the plan is built, the briefing is constructed simultaneously since the MBP enables the user to save 'pre-canned' flyouts and map views configured to illustrate operational issues.

The MBP 'Outline Plan' can be analysed using various displays to show:

- Summary of Missions In Time (Day/Night Capable Missions and Sunrise/Sunset Times), see Figure 3-4.
- Summary of Air Unit Usage.
- Apportionment of resources to mission roles.
- Battlegram of sorties per mission role over time, see Figure 3-5.

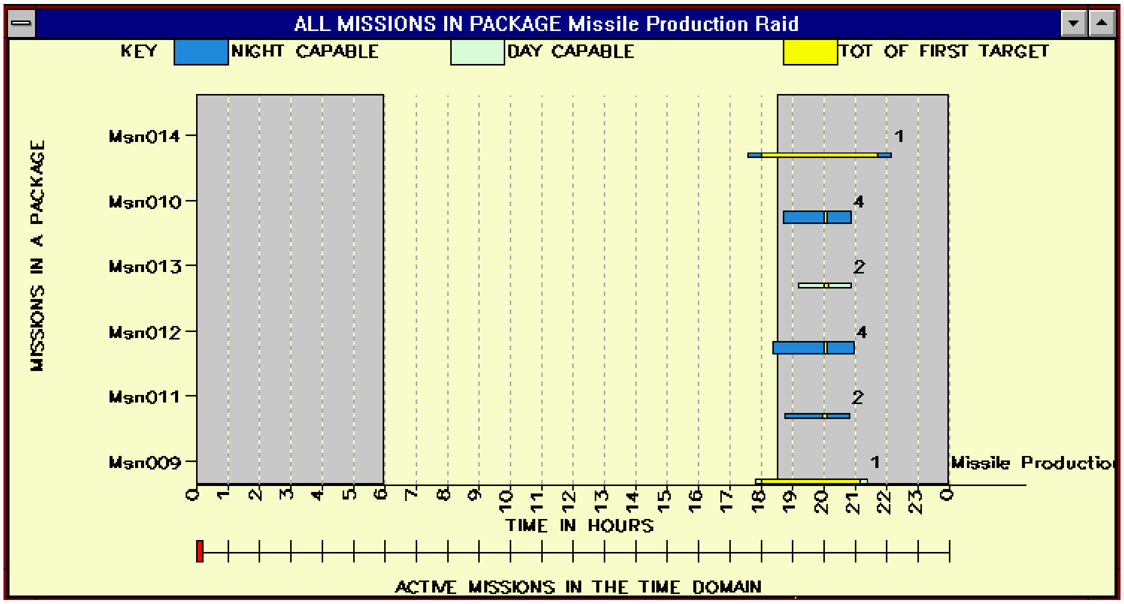

Figure 3-4: Summary of Missions in Time.

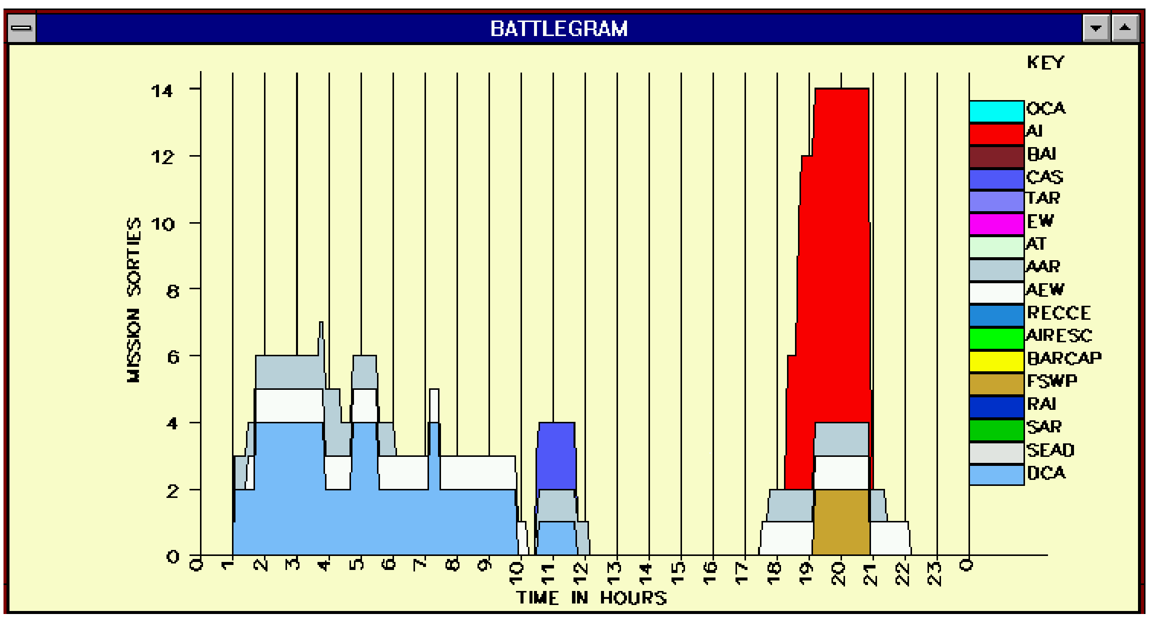

Figure 3-5: Battlegram.

Furthermore, the MBP system is capable of performing the air campaign planning task totally independently of other systems. Where there are other planning systems in use for an operation, the MBP is capable of importing data (such as theatre setup and target information) from them. The MBP has an interface to, for example, the following systems; the US-developed Contingency Theatre Automated Planning System (CTAPS), and the NATO-developed Interim CAOC Capability (ICC).

A certain element of 'what if' or 'contingency plan' development is supported by the MBP system. For testing the derived plan against a perceived enemy threat an interface has been developed to a wargaming system, also produced at QinetiQ (formerly known as the Defence Evaluation Research Agency Malvern), so that the MBP can support the White Team. This wargame system is designed to stimulate the operational thinking in the warfighter and is called the Stimulator.

3.2 SYSTEM EVALUATION

The MBP has been evaluated both by military and non-military users in order to assess its ease of use in an operational environment by different levels of users. The assessment was both structured and intuitive to assess the effectiveness of the MBP's visualisation interface.

3.2.1 Non-Military User Evaluation

The system was evaluated by a civilian. The following is a summary of his report [2]:

> *The MBP is an extremely well written and user-friendly program. As one of those people who dislike reading instruction manuals, I found it very easy to jump straight in and experiment with the wide but logical array of functions. The program is so intuitive; in fact, that I believe it is definitely possible to master it in less than one hour. This would be a great advantage in that training times and hence costs would be reduced, a particularly important factor in the field.*
>
> *The icons for buttons and map objects have been assigned logically, and it is clear that functionality was an overriding consideration in the design of the whole program. Whereas many programs would have opted for fancy, multi-coloured graphics, which one would need to guess the meanings of, the MBP goes for the simplest and most obvious icons. Most option menus have been constructed in similar fashion; as a result, once you have learnt to navigate around one menu, you do not need to learn anything else to use all others. This fact renders much of the manual unnecessary, through no fault of its own, because it contains very clear explanations of all aspects of the program.*
>
> *The range of available functions is sufficient and intelligently chosen. I felt that "Distance between points" and "ATO Previewer" were two especially handy features, the latter allowing (rather entertaining!) animations to be displayed with a click of a button. The overtasking alert was invaluable as well. It would be very interesting to evaluate the program again after the planned introduction of 3D terrain and land heights. This would undoubtedly bring a new level of realism into the plans and scenarios.*
>
> *One small bug that I would like to see ironed out in future versions is one that occurs when the map outline feature is used in conjunction with the ATO Previewer. It causes the display to switch between dark and light for no apparent reason. However, the program is otherwise remarkably stable and bug-free. The Master Battle Planner is an impressive piece of software and a joy to use.*

3.3 MILITARY USER EVALUATION

In 1997 the Command and Control Battle Management Battlelab (C2BMB) identified the Air Tasking Order (ATO) visualisation and assessment capability as an innovation that would enhance the effectiveness and efficiency of the Air Operations Center's (AOC) ATO production process [3]. To measure the potential of this capability, the C2BMB assessed the capability of the technology available at the time to graphically display and assess ATO information.

The C2BMB assessment team evaluated the operational utility of ATO visualisation and assessment capabilities to the functionality of the AOC. Measures of merits (MoM) were developed to determine if an ATO visualisation and assessment capability was suitable for integration into an AOC. The criticality of each MoM was rated as 'essential' or 'desired'.

To effectively define operational requirements, the C2BMB identified eight fielded systems that provided some capability to graphically display, in two or three dimensions, elements of ATO information. The MBP was one of the eight candidate systems. All eight systems were designed to operate in a standalone manner and are fielded, or would be fielded, with US and NATO air forces.

All eight systems were examined during the Exercise Warrior Flag 97 which took place on 19th – 29th August 1997, at the USAFBTS, Hurlburt Field, Florida.

Subject Matter Experts (SMEs) from Air Mobility Command (AMC), Air Force Special Operations Command (AFSOC), Public Air Forces (PACAF), US Air Forces Europe (USAFE), and 9 Air Force observed the capabilities and features of the eight systems during the assessment. The SMEs assessed the criticality of the capabilities (i.e. essential, desired, or non-essential) and the potential utility of the capabilities in the Strategy, Combat Plans, Combat Operations, and Air Mobility Divisions.

3.3.1 Evaluation Process

The assessment team provided the exercise ATO, ACO, unit lay down, and enemy orders of battle information for parsing, display, and analysis. A single set of static data was provided to the ATO visualisation systems from a CTAPS local area network (LAN) that was isolated from the exercise CTAPS network (but dependent on its data). Additional data (e.g. weather) was provided, as requested.

SMEs observed the capabilities and features of all the systems during the assessment. The SMEs then assessed the criticality of the capabilities (i.e. essential, desired, or non-essential) and the potential utility of the capabilities in the Strategy, Combat Plans, Combat Operations, and Air Mobility Divisions. This assessment was supported by Electronic Systems Center (ESC), Air Force Operational Test and Evaluation Center (AFOTEC), 505 CCEG, and Air Force Development Center (AFDTC).

The MBP was set up, operated, and maintained during the exercise. The MBP was demonstrated to the SMEs twice during the course of the evaluation. During the exercise a user manual was available. In the first iteration, the MBP was demonstrated to the SMEs and the SMEs were allowed minimal hands-on interaction. During the second iteration, the SMEs became more familiar with MBP and assistance was provided when required.

3.3.2 Results Data Management Protocol

Questionnaires were the primary data collection tools. The SMEs completed background questionnaires prior to the evaluation. The SMEs also completed assessment questionnaires prior to the evaluation to document preconceptions of the ATO visualisation and assessment capability. At the end of the evaluation, the SMEs completed the same questionnaire to record their opinions after observing two demonstrations of the MBP.

Questionnaire responses were collected and collated during the assessment. SME comments were also recorded during the interviews conducted after each demonstration. At the end of the assessment, the MBP technical expert completed assessment questionnaires to enhance the assessment team's understanding of their systems. MBP system information was also used to enhance the technical descriptions of demonstrated features.

The assessment team interviewed each SME during and after the demonstrations and recorded their responses. The SMEs also input assessment worksheets to document their observations about ATO visualisation and assessment capabilities. Discussions were conducted after the first iteration of demonstrations. These discussions were used to determine the operational utility of the demonstrated capabilities and to define a consensus on the importance of the capabilities.

The assessment team collected a large amount of raw data from the MBP demonstration. This data was organized, grouped, and analyzed to determine its significance.

3.3.3 Objectives and Measures of Merit

During the exercise a set of objectives were established to focus the candidate systems' efforts on advancing the Air Force's core competencies. The Air Force's core competencies were:

- Air and space superiority;
- Global attack;
- Rapid global mobility;
- Precision engagement;
- Information superiority; and
- Agile combat support.

The objectives established to advance the competencies were:

- Demonstrate the capability to provide in-process visualisation and assessment for all, or a selected subset, of the ATO.
- Demonstrate and assess the capabilities of fielded and prototype ATO visualisation and assessment systems.
- Demonstrate the capability to build and display an ATO, airspace, threat, and target picture.
- Demonstrate the capability to import an ATO.
- Demonstrate the capability to manipulate an ATO.
- Determine the potential utility of these capabilities in all major areas of the AOC.
- Propose changes to AOC equipment, processes, and training, as applicable.

The following subjects for evaluation (with MoMs) were developed to determine if an ATO visualisation and assessment capability was suitable for integration into the AOC:

- Display an animated fly-out of ATO sorties overlaid on a flight chart populated with geographical, airspace, weather, and intelligence information.
- Display the volume of information consistent with a large air campaign.
- Enhance ATO effectiveness.
- Enhance user productivity with enhanced graphical and display features.

Measures of Capability (MoCs) were the criteria for evaluating the MoMs. The aggregation of SME assessments of the MoCs were used to rate the Battlelab MoMs as essential, desired, or non-essential.

3.3.4 Evaluation Results

The MBP demonstrated capabilities and features that the SMEs determined to be essential to ATO visualisation and assessment. The assessment team concluded that existing technologies fulfilled most of the initial requirements, but neither the MBP nor the other sevens system demonstrated all the essential capabilities.

The ATO visualisation and assessment capabilities demonstrated were determined to have some utility across all the major divisions of the AOC: Strategy, Combat Plans, Combat Operations, and Air Mobility Divisions.

In summary it was found to be essential to:

- Display an animated fly-out of ATO sorties overlaid on a flight chart populated with geographical, airspace, weather, and intelligence information.
- Import and display user-selected geographical, intelligence, weather, and friendly and enemy order of battle information.
- Display volume of information consistent with a large air campaign.
- Process, manage, and display multiple ATOs, ACOs, and changes.
- Manage and display linkages between ATOs, ACOs, and changes.
- Integrate the tool capabilities with TBMCS for access to ATOs, ACOs, and relevant databases (e.g. APS, ADS, CAFMS/FLEX, JPT, CIS, JMEM, and IMOM).
- Import USMTF ATO and ACO files and messages regardless of the source system.
- Provide capabilities to view user-selected textual ATO and ACO data organized by objective, apportionment category, task unit, mission type, package, mission, target, airspace, aircraft type, time, and location.
- Essential. Display user-selected graphical ATO and ACO data on the map/chart displays.
- Provide the capability to view relevant textual ATO and ACO data by clicking on graphical display items (such as, missions, aircraft, bases, and locations.
- Provide user-configured reports of ATO and ACO data.
- Provide graphical charts of ATO and ACO data to include mission and package flow and sequencing, planning status, apportionment, target allocations, unit, base, and aircraft allocations, and airspace.
- Enhance user productivity with enhanced graphical and display features.
- Provide a DII COE compliant user interface.

3.4 CONCLUSIONS

Overall, exercises and operational evaluations, in the UK and as part of NATO and coalition exercises, have shown that the MBP is an invaluable aid to improving situation awareness and reducing plan development time. When projected onto a touch sensitive display, the MBP is a perfect focus for teamwork.

An additional benefit is that the MBP system provides an excellent means of briefing the plan to the Commander, saving time and reducing the IT support required.

3.5 REFERENCES

[1] Ellis, M., User Manual for the Master Battle Planner, DERA/LS/(SEC-M)/MBP/SUM/2.10, September, 1997.

[2] Ho, G., Master Battler Planner Evaluation, Internal report, July 2000.

[3] Colonel Carpenter, M., Battlelab Commander, Kenney Battlelab Initiative After Initiative Report: ATO Visualization and Assessment, October 10, 1997.

Chapter 4 – A SUMMARY OF RESEARCH CONDUCTED ON THE COMMAND VISUALISATION TESTBED

J.G. Hollands and Jocelyn M. Keillor
Defence R&D Canada – Toronto, Canada

4.1 INTRODUCTION

4.1.1 Background

NATO Research Task Group IST-021/RTG-007 interpreted visualisation as a human activity supported by technology. Visualisation is therefore a means by which humans make sense of complex data. The RTG considered visualisation technologies, including display devices and techniques in relation to how they help humans to perform their tasks effectively. The RTG emphasized the human use of the computational subsystem in ensuring that the right information is available in the form and at the time needed.

An Action Item was assigned to the IST-021/RTG-007 Canadian Point-of-Contact, requesting a section for the Final Report on progress with the Defence R&D Canada – Toronto (DRDC Toronto) Command Visualisation Testbed (CVT). As originally conceived, existing testbed facilities at DRDC Toronto were to be used to examine several prototype command-and-control related systems involving advanced visualisation concepts developed in the various NATO nations. These included: NORCCIS (Norway), Master Battle Planner (MBP) (UK), and SICCE (Portugal). However, despite repeated attempts it was not possible to obtain working versions of any of these systems. Thus, evaluation of the various national systems could not be performed.

Nonetheless, work relevant to IST-021/RTG-007 was conducted in the DRDC Toronto visualisation laboratory, funded under the auspices of the DRDC Technology Investment Fund (TIF). This work involved the empirical evaluation of a wide range of visualisation display concepts (detailed below). This chapter will describe that work. Implicit in the approach is the assumption that the effectiveness of visualisation display concepts is a question best assessed empirically – that is, by having human observers attempt to make judgments with different display arrangements, and collecting performance data from them as they do so. While this approach is not without its shortcomings, it offers an understanding of the factors that affect human judgment with visualisation systems.

4.1.1.1 Testbed History

Command and control systems need a single, intuitive, common tactical picture that displays the necessary information to increase battlespace awareness [1]. However, tactical displays provide poor support for battlespace visualisation [2] and for the cognitive tasks of commanders [1]. In response, the Canadian Department of National Defence (DND) has identified leading-edge technologies for information management [3] as a focus for DRDC, and information and knowledge management for decision making in a complex environment as one outcome of the research programme.

There is a clear trend towards military systems that handle more and more data, while at the same time there is pressure to reduce the number of personnel in operational systems. Therefore there is a need to produce more effective information management systems. The Naval Command and Control Way Ahead [4] notes that, “Display technologies and interfaces reflect more the technical limitations of the era in which they were designed rather than the true requirements for battlespace visualisation and decision support. Information is presented in a manner that maps inefficiently onto human perception and reasoning processes. … little support is provided for cognitive tasks throughout the decision cycle, including planning, intelligence preparation of the battlespace, course of action development and selection, mission rehearsal and simulation." The report further lists the development of "advanced approaches to battlespace visualisation techniques for the naval tactical environment" as an R&D goal.

Further, to augment communication across various command clusters, commanders have a need for a Common Operational Picture (COP). The COP has been defined as the integrated capability to receive, correlate and display heterogeneous sources of information in order to provide a consistent view of the battlespace. A fundamental problem with the communication process for collaborative working is that of the co-ordination of views on an information space, including geographic terrain.

A common finding in research on graphical displays is that the effectiveness of various graphical formats is task dependent [5,6]. Thus, for the range of tasks involved in command and control situations, an effective display must allow users to transition smoothly from one format to another. Typically, proposals for future battlespace displays show a map or chart background with relevant data superimposed in appropriate geographic locations, either graphically (line, bar graphs, timelines) or in tabular form. Some displays show the scene in geometric perspective (God's eye view) and use techniques and algorithms developed for scientific and information visualisation. However, the underlying human factors of using such displays are not well understood and their effective contribution to military command tasks has not been shown.

4.1.2 Purpose and Scientific/Technical Objectives

In response to these client needs, research scientists in the Human-Computer Interaction (HCI) Group at DRDC Toronto proposed the development of the CVT, so that experiments measuring human performance in representative tasks could be conducted. CVT therefore provides the capability to investigate whether proposed visualisation algorithms, constructs, and display concepts are consistent with human perception and cognition and whether they improve command decision making.

Specific scientific objectives included gaining a better understanding of techniques that reduce disorientation when an observer makes quick shifts from global to local views of the battlespace, developing methods for reducing perceptual bias in magnitude judgments of graphical battlespace elements, examining sets of mental operations used in command visualisation in order to maximize performance efficiency, and developing methods for establishing how much information is available from a display "at a glance".

4.1.3 Approach

The CVT includes state-of-the-art graphics workstations, display hardware, three-dimensional (3D) graphics software, and two eyetrackers. The facility was built with the aim of initiating a multi-experiment research program to determine how to relate human perception and reasoning processes to the elements of a command battlespace display. Using a range of military tasks, the collaborators investigated four general themes.

1) *Frame of reference and visual momentum.* The utility of two-dimensional (2D) and 3D displays to depict terrain information has been extensively investigated and the results indicate that the effectiveness of such displays depends on the judgment task. Since the task changes with context, the commander may need a variety of displays to accomplish various ends, leading to the need to switch displays periodically as tasks shift. Problems include disorientation and the need to mentally perform spatial transformations when transitioning from one format to another. Various visual

momentum [7] techniques are available to help commanders transition between two-dimensional (2D) and three-dimensional (3D) displays. This includes techniques such as smooth rotation and tethering.

2) *Perceptual bias and reference points.* Human judgments of the geometric volumes and areas that are commonly used to depict quantitative values in 3D data representations in statistical graphs and maps are biased [8]. Previous work has shown that reference points can reduce judgement error in graphical displays such as those used in command visualisation systems. Two further questions were investigated:

   a) The effect of response method on perceptual bias, and

   b) Bias engendered by the use of perspective rendering in 3D displays.

3) *Modeling mental operations.* Follette [9] proposed that two factors affect quantitative judgments with graphs:

   1) The number of operations necessary; and

   2) The effectiveness of the perceptual features used as input for the operations.

   Empirical work was conducted to determine the effects of these factors on error in quantitative judgments made with graphical displays.

4) *Visual attention and visual span.* Even when a tactical display accurately depicts all relevant data, the human observer may not attend to all displayed elements. The effectiveness of different symbologies to provide relevant tactical information "at a glance" was demonstrated using a change-blindness paradigm. Using an eyetracker, a gaze-contingent display can be constructed that depicts only that region of the display upon which the observer fixates. By varying the size of that region using a staircase procedure, the result indicate how much display information can be attended "at a glance" by a human observer.

More specifically, each research theme provides an improved understanding of how information processing in the command visualisation context can be improved. Results from Theme 1 (Frame of Reference and Visual Momentum) identified effective techniques for transitioning from specific (immersed) to general (world-centred) reference frames, or vice versa. Theme 2 (Perceptual Bias and Reference Points) results provided improved understanding of the selection of display variables to code different kinds of battlespace and command information, and of the utility of reference points to help reduce judgment error for dimensions used in 2D maps and 3D terrain imagery (e.g. area, volume). Theme 3 results (Modeling Mental Operations) provided better understanding of the nature of information processing by testing a framework of mental operations in the command context. Theme 4 (Visual Attention and Visual Span) demonstrated methods for increasing the speed and detectability of tactical display targets. In sum, the results obtained from the proposed work will generally enhance the efficiency of information processing using command visualisation systems.

4.1.4 Delivery Methods/Preliminary Data

Empirical data collection was conducted with the CVT at DRDC Toronto. Experiments used the subject population available to DRDC Toronto (i.e. military and civilian staff, student research assistants). Performance was quantified through the use of time and accuracy measures. Some subjective measures were also recorded (e.g. ability, workload, and preference data, but are not reported here for brevity). Results from the research have been demonstrated to the client group for potential incorporation into future programs, to researchers at other DRDC Centres, and to scientific and military personnel from various NATO and TTCP nations. Prototypes of visualisation techniques and a general purpose 3D graphics engine were developed. The results have been disseminated through presentations at peer-reviewed conferences and workshops, and submitted for publication in scientific journals. This chapter provides a summary of the work, and treats the results from each research theme in turn.

4.2 THEME 1: FRAME OF REFERENCE AND VISUAL MOMENTUM

4.2.1 Theoretical Background

The intent of the COP is to provide a shared understanding of the battlespace to improve responsiveness and provide decision dominance. Visualisation technology offers a means to establish the COP and should help the commander transition across strategic, operational and tactical levels. The ultimate aim is to obtain an integrated visualisation environment where commander and staff can gain a shared understanding of the changing battlefield situation.

Current technology provides a means to display tremendous quantities of data to the human commander. Geospatial data, sensor data, network data – electronic and human, socio-political data, data on troop status, materiel, data from news media, and so on. Further, the nature of the future operational environment requires a high degree of flexibility and adaptability, due to factors such as: asymmetric threat, enlarged areas of operation; non-contiguous and non-linear operations; requirements for a three-block war, use of complex terrain, and effects-based operations within a maneuverist approach [10].

To address this state of affairs, it is argued that the future commander can only be successful if the command team functions collaboratively. A rigidly hierarchical command structure is too slow and inflexible to respond in a timely manner to the nature of the asymmetric threat and resulting warfighting characteristics. There is a need therefore, to provide an integrated command and control system that supports collaborative working. This includes the design of computer software and displays to facilitate the collaborative working concept.

While co-ordinating information is shared across different echelons, commands, environments, government departments, and nations, the information available varies across such organizations, and is often represented and portrayed in different ways. A fundamental problem with the communication process for collaborative working is that of the co-ordination of views on an information/knowledge space. For instance, if a shared geospatial awareness is required – the platoon commander with troops on the ground looking at a group of buildings versus the company commander examining a 2D plan view (aerial photographs, maps) of the same urban terrain – it can be difficult for one commander to communicate to the other. In one view, task-relevant information may be visible: in the other, invisible. Is the company commander aware of what is visible to the platoon commander? What is left and right in the forward field of view (FFOV) may be reversed with the map depending on orientation. If a shared understanding of network access data is required (e.g. a complex set of intrusion detection data) two analysts in different locations may have access to only limited views on the data, but need a common representational format (e.g. a dynamic 3D graph) to communicate.

Hollands, Lamb, and Keillor [11] argued that such problems are related to the frame of reference concept [12,13,14] and present a framework for that concept that classifies factors that improve or degrade performance when co-ordinating information across views of spatial data. They also considered similar display concepts from information visualisation in human-computer interfaces (depicting file structures, networks, the web, windows and other interface elements; see [14]) and from display design guidelines from other domains (e.g. process control, medical imaging), and note the fundamental similarities. In particular the relevant literatures underline a recurring need for depicting both global context and local content, which leads to the need for multiple displays. Further, since different viewpoints have various advantages and disadvantages for various operational contexts, there is a need to provide methods for improving visual momentum across displays.

Hollands et al. [11] note how such methods can be split into two basic types: those that allow both views to be shown simultaneously (compromise displays) and those that ease the abrupt shift between global and local views by showing the mapping between display elements in different views (transition displays). Hollands et al. [11] also presented a preliminary taxonomy that lists the factors that distinguish egocentric from exocentric displays. This includes the distance and angle of elevation of the viewpoint with respect to a visual scene (Level I), the distance and angle of elevation with respect to a particular object of interest within the scene (Level II), distance and viewing angles with respect to the motion of an object or objects of interest within a scene (Level III), and the distance and angular compatibility between the viewpoint and the motion of a controlled object within the scene (Level IV). In the Level IV situation, the control order dynamics and the rigidity of the link between the viewpoint and controlled object are also of concern.

The Hollands et al. [11] taxonomy can be used to predict performance: as the number of shifts in reference frame increases, increased time and error are predicted. This is because each shift is essentially a transformation in reference frame; each transformation requires time and there is some likelihood of error if the computation required is not performed accurately [15,16,17]. This would be true for communication of shared understanding among collaborating workers, or for shifts in viewpoint over time for a single individual. The classification should lead to the appropriate use of display techniques to help the commander minimize mental effort, maintain good situation awareness, and improve communication across different levels of command, coalition partners, and public or private agencies.

We can also consider how multiple sensors can provide a new interface for controlling and monitoring platforms and their surrounds. Sensor data streamed from a moving vehicle integrated with data from other nearby sensors will provide remotely operated vehicle (ROV) operators and their commanders with reconstructed 2D and 3D visual representations of complex terrain. Virtual or augmented reality visualisations allow the use of viewpoints different from that "out the window". Designers are then faced with the choice of optimal viewpoint parameters that maximize human performance.

It is widely accepted that the nature of the task dictates the best viewpoint on geographic terrain. Tasks involving shape understanding are best performed with 3D viewpoints because all three dimensions are integrated into one representation. 2D viewpoints are best for precise tasks judging relative position, due to the distortions associated with 3D viewpoints [18]. For navigation and wayfinding, local guidance is best performed from an egocentric perspective, while global spatial awareness tasks should be carried out with an exocentric, fixed viewpoint showing most or all of the terrain [14].

The implication, then, is to provide the appropriate viewpoint for the relevant task. However, as we have noted, switching between multiple displays is disorientating and leads to the need for spatial transformation. Each added transformation requires extra processing time and increased likelihood of error. We describe here two approaches we have taken to this problem. One involves the use of a tethered display, which serves as a compromise display incorporating both egocentric and exocentric elements. The other examines the utility of smooth rotations between display formats, serving as a transition display. Both methods, we argue, provide good visual momentum [7] between egocentric and exocentric display formats.

4.2.2 Tethering

The tethered viewpoint [19,20] is commonly used in computer gaming and couples the viewpoint to the position and orientation of a moving object (or *avatar*). The viewpoint is typically higher than the avatar and behind it, showing more of the terrain than would be seen from the avatar's viewpoint. In this sense, it provides visual momentum by providing a view that incorporates egocentric and exocentric qualities. There is some evidence in support of the tethered concept: Wang and Milgram [21] found that a tethered display produced better performance than an egocentric display for aerial navigation, and Wickens and Prevett [20] found advantages with a tether-like display (versus egocentric) for spatial awareness.

However, a rigid tether violates the principle of motion compatibility. As the operator directs the avatar to the right, the visual scene moves to the left. The rigid tether also behaves like a compensatory tracking system [19]. Modelling the tether as a dynamic mass spring damper system creates a display incorporating both compensatory and pursuit tracking attributes and may reduce the motion compatibility problem. However, Wang and Milgram [21] did not find any advantage to the dynamic tether over a rigid tether for six-degrees-of-freedom (DOF) aerial navigation.

Ground-based navigation differs from aerial navigation in that the operator is only controlling one rotational DOF (yaw). (Although pitch and roll vary when driving, they are typically not under the direct control of the operator). Yu and Andre [22] used an arcade driving simulator to evaluate four different viewpoints. A tethered viewpoint slightly removed from the vehicle provided the best navigation and situational awareness. However, this driving task differs in several fundamental ways from off-road navigation in complex terrain.

Using the DRDC Toronto CVT, Lamb and Hollands [23] examined the effect of viewpoint on controlling a vehicle navigating complex terrain, and on concurrent and subsequent spatial awareness. Five viewpoints were used (see Figure 4-1): egocentric, rigid tether, dynamic tether, 3D exocentric (perspective view on terrain, but viewpoint does not change with vehicular position), and a 2D exocentric map (God's eye view). Lamb and Hollands predicted that navigational performance would be better with the egocentric than either exocentric view because the egocentric display provides well-learned egomotion cues commonly available as one navigates through an environment, either in vehicle or on foot [14]. Lamb and Hollands also predicted that the tethered display would be more effective than the exocentric displays for navigation, and as effective as the egocentric display. Finally, they predicted that navigational performance would be better with dynamic than rigid tethering, reducing the motion compatibility problem and allowing pursuit tracking.

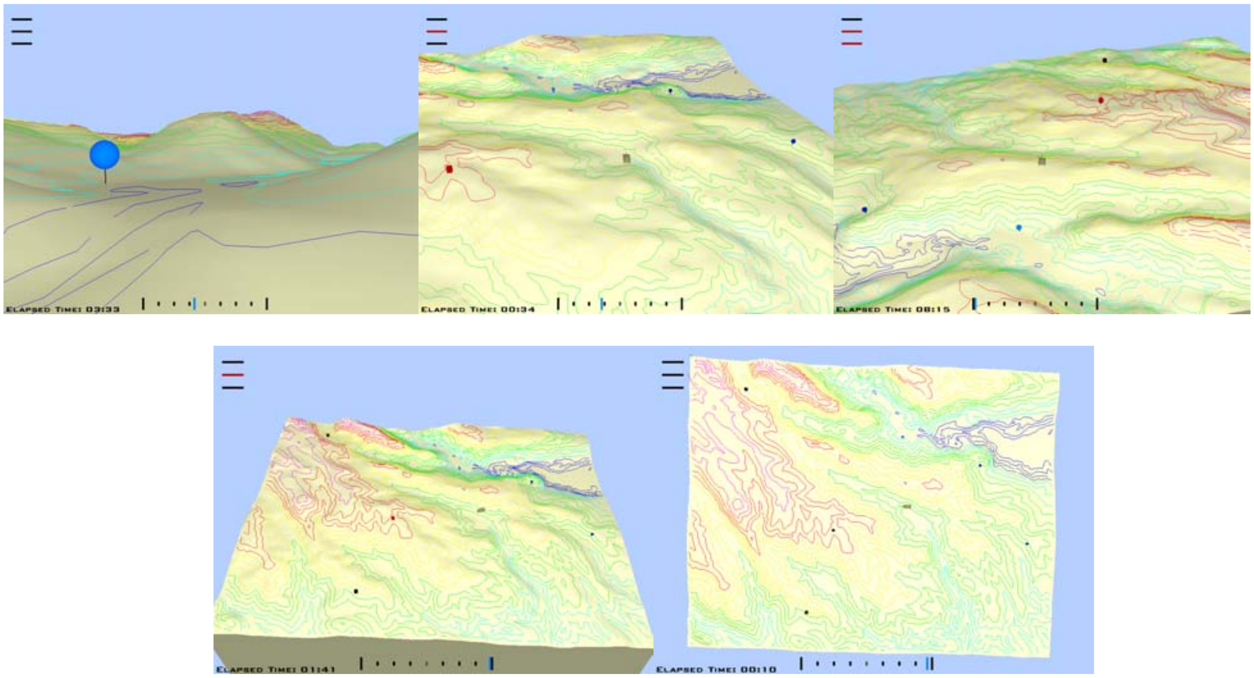

Figure 4-1: Example Views of each Display Type. Top row (left to right): egocentric, dynamic tether, rigid tether. Bottom row: 3D exocentric, 2D exocentric.

In the Lamb and Hollands [23] study, participants were instructed to avoid being seen by enemy units while navigating the simulated vehicle between waypoints. This was done to assess spatial awareness during navigation. The time that the tank was seen from at least one of the enemy positions was recorded. After each trial, observers had to choose the terrain just navigated from a set of distractor terrains. Given that both of these spatial awareness tasks require a sense of the global characteristics of the terrain, Lamb and Hollands predicted the opposite order of effectiveness: exocentric worse than egocentric, with the tethered display as effective as the exocentric display. They did not predict any effect for dynamic (versus rigid) tethering for these tasks.

Consistent with predictions, Lamb and Hollands [23] found that the egocentric display was more effective than exocentric displays (2D or 3D) for navigation, and the exocentric displays were more effective than egocentric for spatial awareness, both for time seen during navigation and the recognition task. The tethered displays generally produced intermediate results. For navigation, they were less effective than the egocentric display and roughly equivalent to the exocentric displays. For spatial awareness recognition the tethered displays were more effective than the egocentric display, but less effective than the exocentric displays.

More importantly, the tethered displays were the most effective displays for spatial awareness for minimizing time seen. Not only were the tethered displays more effective than the egocentric display, they were also more effective than the exocentric displays. Use of the tether minimized the time during which the participant's avatar was visible to enemy positions. There was no effect of tether dynamics in the navigation task.

In summary, the tethered display was useful in spatial awareness involving knowledge of locations of interest with respect to one's own position while navigating. In this sense, perhaps the Lamb and Hollands [23] results identify a navigation task whose performance is dissociated from conventional exocentric spatial awareness and egomotion. The tethered display may be the most effective display for this type of egocentric spatial awareness task.

4.2.3 Visual Momentum and Smooth Rotation

Background. In a series of experiments using the DRDC Toronto CVT, Hollands and co-workers [24-26] have examined the utility of *visual momentum* for the depiction of geographic terrain, a topic that has received relatively little empirical attention despite its clear importance. The concept was first defined as such by Hochberg and Brooks [27], who described techniques used in film to help an audience maintain spatial understanding of a scene across discrete film cuts. Woods [7] extended the visual momentum concept to user-computer interaction, and defined it in that context as the user's ability to extract and integrate data from multiple consecutive display windows. More recently, Wickens and Hollands [14] summarized four basic guidelines for improving visual momentum: consistent representations; graceful transitions; highlighted anchors; and world maps. Various specific approaches to improving visual momentum have been proposed or implemented [28-31], and examined empirically [29,32-36].

Hollands and co-workers were particularly interested in the problem of the depiction of geographic terrain for command and control. This is an important component of battlespace visualisation systems [37]. As described earlier, there is benefit to providing multiple views on terrain, and therefore both 2D and 3D display formats should be made available to the observer. In the command and control context an observer needs to switch tasks frequently while viewing geospatial information, leading to spatial disorientation and the need for spatial transformation. A gradual transition between 2D and 3D perspectives (and vice-versa) incorporating animation of viewpoint during task switching may provide visual momentum and alleviate the problem.

Hollands and co-workers were interested in the question of whether smooth transition aided observers as they switched tasks. To examine this question, they used two tasks developed by St. John et al. [18]. A shape understanding task required the participant to judge whether one ground location was visible from another (*A-See-B* Task), and a relative position task required a judgment of which one of two points was of higher altitude (*A-High-B* Task). St. John et al. found that the A-See-B task was performed better with a 3D display, whereas the A-High-B Task was performed better with a 2D topographic map.

Hollands and co-workers had participants switch tasks across trials to determine whether knowledge of terrain obtained when performing one task affected performance on a different task on a subsequent trial. In a *continuous transition* condition, the display rotated in depth from one display format to the other. The design of the control condition (discrete transition) varied across experiments. In Experiment 1 [24], a blank screen was shown for a duration equal to that used for the continuous transition. Participants were immediately shown the alternate display format in Experiment 2 [25]. In Experiment 3 [26] the 3D display was oriented to be aligned with the A-B vector. For the continuous rotation condition, the terrain was rotated in its azimuth prior to the rotation in depth. In the discrete case, the azimuth rotation was in the opposite direction and then the terrain "snapped" to the final orientation to examine the role of preview during transition.

For all experiments, it was predicted that the continuous transition would improve second-trial performance relative to discrete transition. This is because the smooth rotation should provide improved visual momentum between consecutive displays. We provide only a synopsis of the three experiments here; for details of the experimental methods, the reader should consult the original articles.

In general, terrain models were created from Digital Terrain Elevation Data (DTED) using Creator/TerrainPro [38] modelling tools. Each model represented a 13351 m x 11288 m region of the US state of Wyoming.

2D and 3D displays were constructed to resemble those used by St. John et al.[18]. The Vega visual simulation system [39] was used to render each terrain model as a 3D display, and an example is shown in Figure 4-2. The 3D display depicted the terrain model at a viewing angle of 45 degrees with respect to the ground plane. MICRODEM (Microcomputer Digital Elevation Models) [40] was used to create a 2D display with coloured contour lines (see Figure 4-2 for an example). 2D and 3D displays depicting each of the 8 A-B pairs were constructed, resulting in 16 different displays per terrain model. Each location in a pair was represented by a point superimposed on the map labeled A or B.

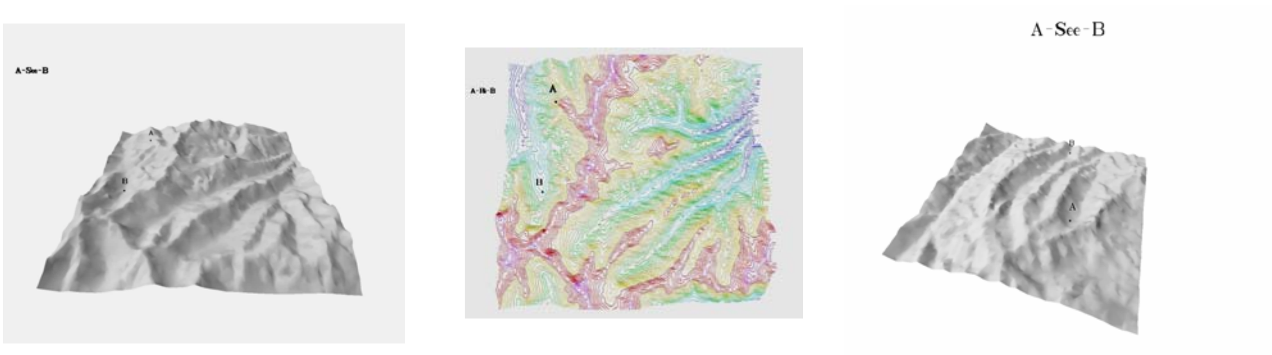

Figure 4-2: Example Views of 3D and 2D Displays Used in the Visual Momentum Experiments [24,25]. The right most 3D display was used in Experiment 3 [26].

Experiment 1. In Experiment 1 [24], participants made judgments about the properties of two points placed on terrain depicted as 2D or 3D displays. They performed the tasks in pairs of trials, switching tasks and displays between trials. On half the trials (continuous transition), the display dynamically rotated in depth from one display format to the other. On the other half (discrete transition), a blank screen was shown for the same duration.

As predicted, the results showed that a continuous transition between the display types improved performance on the trial after transition relative to the discrete condition. Participants were faster and there was a trend toward greater accuracy on the second trial of the pair with the continuous transition. Given the RT advantage without any trade off in accuracy, it appeared that the smooth rotation provided improved visual momentum between consecutive displays.

Experiment 2. Experiment 2 [25] was a replication of Experiment 1 with a key difference: in the discrete transition condition, participants were immediately shown the alternate display format on the second trial of a pair. If the advantage for continuous transition was not obtained in Experiment 2, it would imply that the Experiment 1 results were due to participants forgetting terrain information while the blank screen was shown in the discrete condition. In contrast, if the advantage in Experiment 1 was due to the visual momentum provided by smooth rotation, then the results should still obtain when there is no blank screen in the control condition.

The results of Experiment 2 showed that participants were faster on the second trial of the pair with the continuous transition. As observers switched tasks and display formats, a continuous transition improved performance relative to discrete condition. In the discrete condition of Experiment 2 the second trial was shown immediately following the first, and so this result could not have occurred because participants forgot information gained on the first trial in that condition. Accuracy was generally higher with continuous transition, both before and after transition. This meant that there was no evidence of a speed-accuracy trade off with respect to the transition effect – accuracy was higher after the continuous transition.

Experiment 3. Because the terrain was shown during the smooth transition in the continuous condition, participants in Experiments 1 and 2 may have used this preview to anticipate the second trial of the pair. Was the source of the obtained advantage for the continuous condition due to preview or improved momentum? In Experiment 3 [26] this question was tested by showing the rotating terrain in the control condition for as long as in the continuous transition condition.

In Experiments 1 and 2, the rotation from 2D to 3D (and vice versa) only occurred in depth. In Experiment 3, the terrain was also rotated in the azimuth so that the viewpoint for the 3D display was aligned with an imaginary line connecting points A and B (see Figure 4-2). This was done to make the 3D display more immersive or egocentric [14], and to provide a method for equalizing the rotation time in continuous and discrete conditions. For the discrete transition condition, azimuth rotation was in the direction opposite to that which occurred in the continuous case. For example, if the azimuth rotation to the A-B vector was 120 degrees counter-clockwise in the continuous condition, then it was 120 degrees clockwise in the discrete condition. Upon reaching this position, the display orientation would immediately switch to the azimuth position aligned with the bottom of the 2D map. Then the horizontal translation occurred, followed by rotation in depth to produce the 2D view. The opposite sequence was used to transition from 2D to 3D.

An advantage for the continuous rotation was observed. There was no evidence for a speed-accuracy trade off for this effect – accuracy was not reduced after continuous transition. We argue that continuous transition to the correct position therefore provided improved visual momentum between displays.

Smooth rotation and visual momentum summary. The primary intent of the visual momentum work was to investigate the effect of continuous transition and the results suggest that the source of improved performance is the uninterrupted flow of terrain views. Presumably, this flow provides improved visual momentum. This is not to suggest that there cannot be an advantage to having preview or that displaying task relevant information during the transition will not aid performance. In many real-world contexts, it is probable that such factors will co-occur. Smooth rotation takes time to portray, and it seems likely that the human observer will use this time to prepare for subsequent task demands. The claim is not that these factors will not have an effect, but rather that their presence is not required to produce a performance advantage. The use of dynamic transition is therefore recommended when observers examine multiple views of terrain over time.

4.2.4 Theme 1 Frame of Reference Summary

The research conducted under Theme 1 was concerned with evaluating methods for helping an observer transition between different views of terrain. The results indicate that a tethered view can be a helpful device to monitor one's egocentric position, but still maintain a broader awareness of other relevant locations than is possible with a strictly egocentric view or an exocentric map view. The visual momentum research indicates that smooth rotation provides a flow of terrain views that assists an observer switch between egocentric 3D and exocentric 2D representations. Presumably, these display techniques (tethering and smooth rotation) reduce the amount of spatial transformation required relative to seeing only an egocentric or exocentric view or discrete view switching, respectively.

Future frame of reference work in the CVT will involve conducting empirical studies examining those visual momentum methods useful for switching between different data formats (e.g. tables of materiel for different echelons, vs. maps showing locations vs. graphs and timeline data). These may include linked views between data formats, allowing the user to drag objects from one data format to another [31]. Given the increased interest in collaborative working (ref), we are interested in examining the frame of reference problem and visual momentum techniques as solutions to the communication of information between collaborating partners in a group context.

4.3 THEME 2: PERCEPTUAL BIAS AND REFERENCE POINTS

Human judgments of the geometric volumes and areas that are commonly used to depict quantitative values in 3D data representations used in statistical graphs, maps, and command visualisation systems are biased [8]. However, the use of reference points can reduce judgment error in such judgments. Two research topics were investigated:

a) Bias engendered by the use of perspective rendering in 3D displays; and

b) The effect of response method on perceptual bias.

We first discuss bias in proportion judgments, and then examine each experiment in turn.

4.3.1 Background – Bias in Proportion Judgments

Many real-world tasks require us to estimate the proportion one quantity is of another, larger quantity. For example, we might establish that our gas tank is less than half full by glancing at the fuel gauge while driving. In order for a commander to make an informed decision, it is often necessary for quantities to be compared, leading to the requirement to compute per capita figures from raw data (e.g. number of injured casualties as a function of size of unit). A number of studies have found a consistent pattern of overestimation and underestimation (constant error) in proportion judgments, despite wide variation in task demands and display format: Proportions less than .5 tend to be overestimated, and proportions greater than .5 underestimated [41,42]. Other studies show the pattern cyclically repeating over the range of stimuli tested. For example, Spence and Krizel [43] found a bias pattern whose cycle repeated (over-under, over-under) when subjects judged proportions shown in conventional graphs (e.g. pie chart, divided bar graph). Other studies found a four-cycle pattern (e.g. with angle judgments [44]). Sometimes the pattern reverses (under-over) [45].

Two questions arise when considering these findings. First, what caused the bias to occur? Why, for example, should a small proportion be overestimated, and a large one underestimated? Second, what might account for multiple cycles of bias? Hollands and Dyre [8] proposed a cyclical power model based on Stevens' power law to answer these questions. Stevens' law [46] states that the relationship between perceived magnitude of a stimulus, Φ, and its physical magnitude, Π, is expressed as a power function,

$$\Phi = \alpha\Pi^{\beta}.$$

The coefficient α represents a scaling factor (translating objective to subjective units) and is not that important for current purposes. However, the *Stevens exponent* β indicates the nature of the relationship between physical and perceived magnitude. Response compression occurs when $\beta < 1$; each increase in physical magnitude causes less and less increase in perceived magnitude. (Response expansion, where each increase in physical magnitude causes progressively greater increases in perceived magnitude, occurs when $\beta > 1$). In the magnitude estimation task commonly used by Stevens, the observer is shown a set of stimuli which differ along some physical dimension (e.g. length, area, volume) and is asked to estimate magnitude by assigning a number to each stimulus. Estimates of the exponent of the power function are around 1.0 for length, 0.8 for area, and 0.6 for volume [46]. Reverse patterns [45] can be accounted for using an exponent greater than unity.

Spence [47] proposed a model of proportion judgments based on Stevens' law. Consider two quantities Π and Ω, where $\Pi + \Omega = 1$. The subjective proportion P is computed as

$$P = \alpha\Pi^{\beta} / [\alpha\Pi^{\beta} + \alpha\Omega^{\beta}]$$
$$= \Pi^{\beta} / [\Pi^{\beta} + (1 - \Pi)^{\beta}].$$

This model predicts a one-cycle bias pattern; when $\beta < 1$, small proportions P are overestimated, and large ones underestimated.

Hollands and Dyre [8] proposed a modification of Spence's model [47] to account for multi-cycle bias patterns. Rather than comparing a part to the whole, Hollands and Dyre proposed that the observer can compare the part to intermediate reference points. The cyclical power model proposes that a judgment requires the use of two reference points to define the range of possible responses, one smaller than the to-be-judged proportion, the other larger. That part of quantity P larger than the smaller reference point is compared to the range between the reference points, rescaled as a proportion of the entire stimulus, and added to the smaller reference point. If $R_0, ..., R_n$ is a set of reference points then a proportion may be computed as

$$P(\Pi) = \frac{(\Pi - R_{i-1})^{\beta}}{(\Pi - R_{i-1})^{\beta} + (R_i - \Pi)^{\beta}} \cdot \frac{R_i - R_{i-1}}{R_n} + \frac{R_{i-1}}{R_n} \qquad \text{if } R_{i-1} \le \Pi \le R_i \cdot .$$

The cyclical power model (CPM) predicts one-, two-, and four-cycle bias patterns for P when the number of reference points is varied. When two reference points are used (0, 1) a one-cycle pattern is predicted. When a third reference point is added at .50, a two-cycle pattern results. When reference points are further added at .25 and .75, a four-cycle result occurs. Hollands and Dyre [8] successfully fit the different versions of CPM to one-, two-, and four-cycle data. The experimental data obtained by Erlick [41], Spence and Krizel [43], and Huttenlocher et al. [44] were best fit by one-, two-, and four-cycle versions of the model, respectively.

A key assumption for CPM is that the number of reference points used determines the bias frequency. Placing tick marks in certain locations on a display may affect the choice of reference points. Varying the number of reference points used should affect the bias pattern frequency without affecting the exponent, β. To examine this claim, Hollands and Dyre [48] showed subjects pie charts with tick marks placed at various intervals around the circumference of the pie, on a horizontal response line placed below the pie, or in both locations. The cyclical power model was fit to the data using the procedures described above. Figure 4-3 shows the results with tick marks at halves (0, .50) and at quarters (0, .25, .50, .75, 1), when tick marks were on the graph, response line, or both. The best-fitting version of the cyclical power model is also shown in Figure 4-3. With tick marks at halves the data were best fit by the two-cycle version, as commonly observed with pie charts. In contrast, with tick marks at quarters, a four-cycle version fit best. This occurred only when the tick marks were placed on both graph and response line, or on the response line alone, implying that the selection of reference points may be more related to response production than stimulus perception. Although the frequency of the bias pattern changed, the Stevens' exponent remained constant at about 0.8 as one would expect given that the perceptual continuum did not change. The exponent corresponds well to the exponents obtained from other data involving proportion judgments with pies [47], and is consistent with the typical exponent obtained from magnitude estimates of area. The increase in cyclical frequency generally reduced error, even given a constant Stevens' exponent.

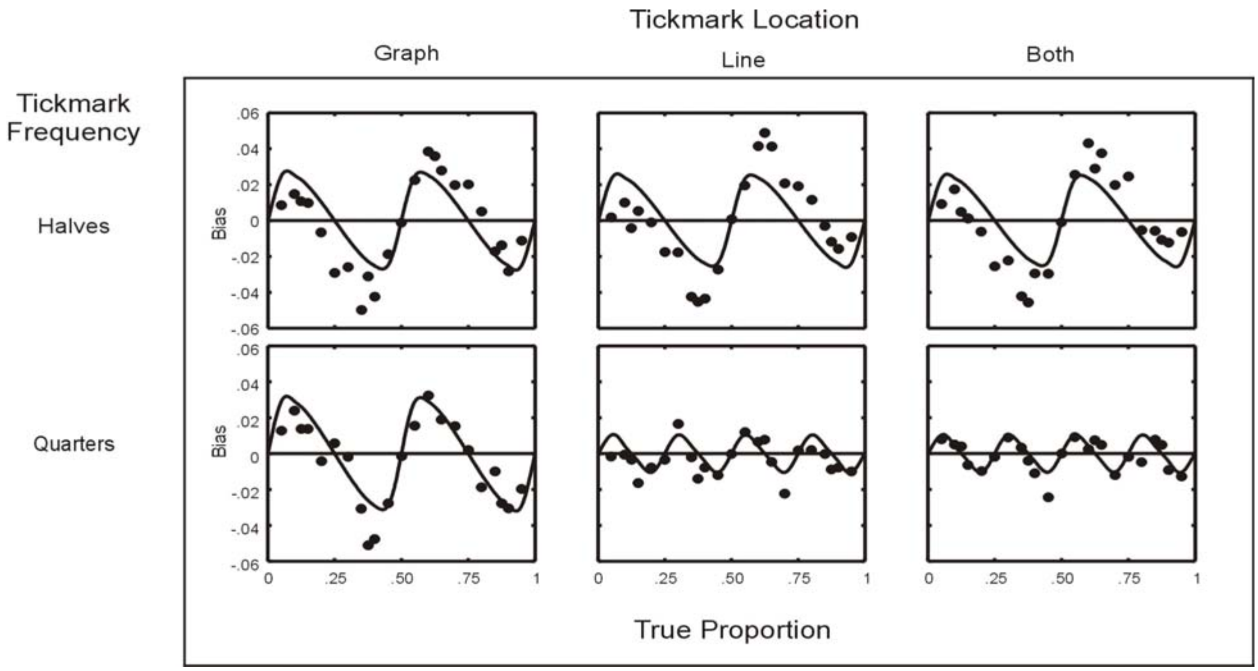

Figure 4-3: Bias (constant error) in Judging Pie Charts as a Function of Tick Mark Location and Frequency [48]. The solid lines represent the predictions of the best fitting version of CPM.

CPM is a general model, and has been shown to account for bias in proportion judgments using a wide range of stimuli [8]. Hence, the model has practical utility for determining how different types of quantitative displays are read. The designer could use the cyclical power model to analyze empirical data comparing various tick mark frequencies and locations to determine if tick marks are being used or not, and also to find the combination producing the smallest bias.

It is important to understand the various types of bias that can affect a perceptual judgment if we are to produce displays that are relatively free of such bias. 3D bar graphs are increasingly used to display complex data of all kinds, including battlespace data [49], and 3D bars are commonly portrayed at different distances from the observer using linear perspective. Does this portrayal at different depths have an effect on the accuracy of proportion judgments? How does it relate to the bias inherent in the judgment of volumes, areas, and lengths discussed above?

Two research topics pertaining to CPM have been investigated using the DRDC Toronto CVT. First, the use of more compatible response methods on the choice of reference points was examined. Second, bias engendered by the use of perspective rendering in 3D displays was investigated.

4.3.1.1 The Effect of Response Method on Perceptual Bias

People make proportion judgments in many situations, and often respond in different ways – by recording a number, by adjusting a control, or by reproducing the position of a display indicator on paper. In the Hollands and Dyre [48] experiment discussed above, participants estimated proportions by dividing a horizontal line into two parts corresponding to the parts of the pie chart. Their results showing that placing tick marks on the response line was the key factor affecting cyclical frequency suggests that something in the nature of the response may affect the choice of reference points utilized. Similar results had been obtained by Taylor [45].

It is a widely accepted human factors design guideline that dial displays are more compatible with rotary controls and linear displays with sliding controls [14,50]. Fitts and Seeger [51] found that such stimulus-response (S-R) compatibility reduced response time and error. Perhaps increased S-R compatibility might also affect the choice of reference points.

Morton and Hollands [52] examined this question by having participants respond using a rotary dial, a horizontal line, or by typing a numeric response. They predicted that greater compatibility between display and response method should increase the number of reference points used and reduce judgment error.

Results showed a two-cycle bias pattern for line and numeric conditions, but a four-cycle pattern for the dial, leading to reduced error for that condition. Response method had no effect on judgment time. Response method did not affect the estimated value of the Stevens exponent (0.83 on average). More compatible S-R relationships between pie chart and dial may have led to the use of higher-frequency reference points, thereby improving judgment accuracy. These results highlight the importance of considering the display-control relationship when attempting to minimize error in proportion judgments.

4.3.2 Biases in Reading 3D Bars

In an attempt to maximize the amount of information available at a glance, modern information visualisation systems make use of display techniques in which 3D graphics are heavily relied upon. In such systems 3D graphs are increasingly used to display complex data of all kinds, and 3D bars are commonly portrayed at different distances from the observer using linear perspective. For example, this occurs when bars are placed at different locations on a map or terrain surface viewed at an oblique angle. It also occurs with 3D bar graphs or histograms where two axes code variables and bar height is used to represent a third variable. However, the use of a 3D representation introduces the potential for error in comparisons of the elements portrayed at different depths.

To give the impression of depth, linear perspective is often used, which results in more distant objects being portrayed at smaller sizes. This mimics the monocular information available when we examine a real 3D scene (e.g. looking out the window with one eye) or when we look at any 2D surface depicting a 3D scene (e.g. a photograph). However, there may be insufficient depth information in an artificial 3D scene to specify object distance and therefore the observer's size-distance scaling may be inaccurate. For 3D displays, depth cues can be added to allow the sizes of objects at different distances to be judged more accurately. Hollands, Parker, and Morton [53] therefore posed the following experimental question: Does size-distance scaling allow accurate judgments of portrayed size in 3D bar graphs given the depth cues typically available (linear perspective, texture, relative height and occlusion)?

Participants were shown two 3D bars placed on a grid, and estimated the proportion the smaller bar at front was of the larger bar, as shown in Figure 4-4. Bars were co-located, placed side by side (near adjacent), or the small bar occluded the large (near or far occluded, respectively).

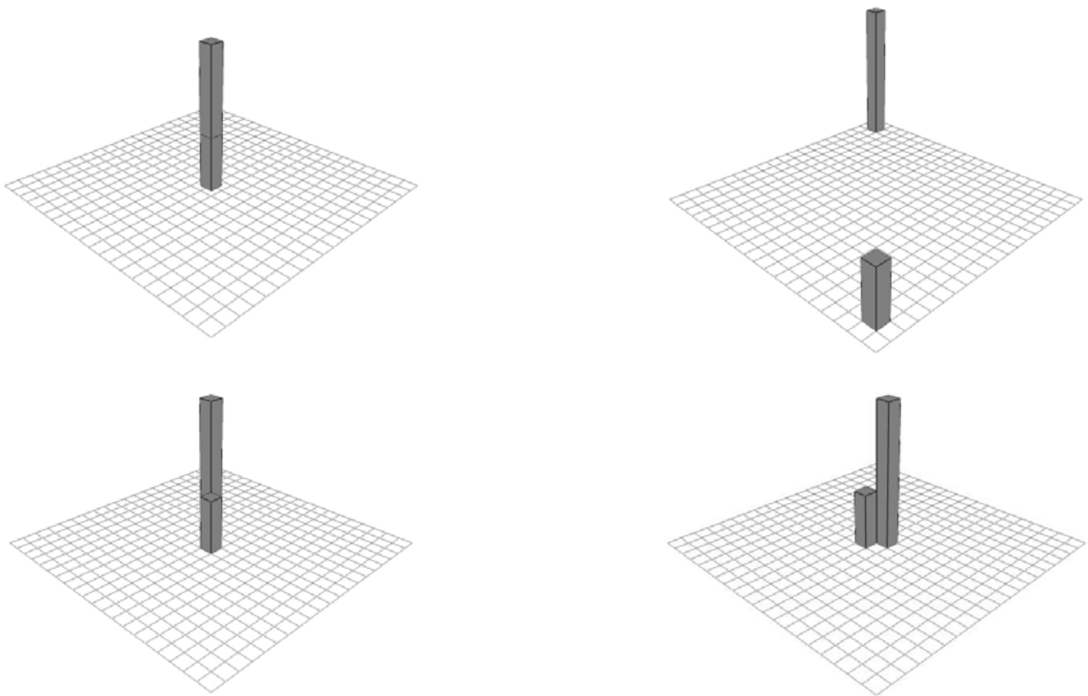

Figure 4-4: Examples of Displays Used in [53] (top: co-located, far-occluded; bottom: near-adjacent, near occluded).

Three types of results are possible if size distance scaling does not hold. If observers judge retinal size (visual angle) then this would lead to general overestimation of the proportion or a positive bias (since the "true" size of the rear bar will be underestimated). If observers perceive the far bar as farther than actually depicted, an overconstancy would result leading to overestimation of the far bar, underestimation of the proportion, and a negative bias. A third possibility is that size estimates simply become less stable, with both overconstancy and underconstancy (either through judgments of retinal size or otherwise) occurring. In this case, there may be no change in net bias, but an increased variability in the scores.

The results showed that judgment error was greater for the far occluded condition and there was greater variability in bias scores across observers. Thus, size estimates in 3D bar graphs were less stable when the bars were placed farther apart, with both overconstancy and underconstancy occurring. Presumably, the available depth information did not accurately specify the relative distances of the two objects, and this poor specification led to imprecise size estimates. In addition, cyclical bias was observed in participants' judgments, with all participants showing best fit with a one- or two-cycle version of CPM.

To account for the observed constant error, Hollands et al. [53] augmented CPM to include an independent bias parameter γ representing size-distance scaling. According to the model, these two types of bias are independently responsible for the observed judgment error. The model was fit to the data with good results. In particular, the absolute value of γ was greater in the far-occluded condition. Inclusion of γ increased R^2 for far- and near-occluded conditions only. In contrast, bar location did not affect cyclical bias.

An implication for graph design is that bars should be portrayed at similar depths in 3D bar graphs if accurate judgments are necessary, or that 2D graphs (where there is no variation in depth) should be considered as an alternative. 3D bars may be augmented by graduated axes scaled to various depths; research is planned to examine the effect of these augmentations on size-distance scaling.

4.3.3 Summary

We draw four conclusions about bias in proportion judgments with commonplace graphical stimuli found in battlespace visualisation systems:

1) Cyclical bias can be expected;
2) The effects of this bias on judgment error can be reduced by adding reference points;
3) More reference points can be added by increasing the frequency of tick marks; and
4) More reference points can be added by increasing the S-R compatibility of the response method and the display arrangement.

4.4 THEME 3: MODELING MENTAL OPERATIONS

4.4.1 Background and Framework

Many studies in graphical perception have shown that the effectiveness of different types of graphs depends on the task [14,54]. For example, Follette and Hollands [55] showed participants bar graphs depicting two quantities A and B. The two bars were either side by side or one was placed above the other (stacked bar). They found that performance on a part-to-whole [A/(A+B)] proportion judgment was better with the stacked bar arrangement, whereas performance on a part-to-part proportion judgment (A/B) was performed better with the side-by-side bar arrangement.

Why does task dependency occur? Presumably, something differs in terms of cognitive processing. Hollands and Spence [56] contend that there must be a difference in mental operations to account for task dependent results. Similar arguments have been made by other researchers [57,58] and models of such processing have been proposed. In general, the models assume that operations are executed in a sequence of steps, and that the time required for a user to complete a task with a graph is linearly related to the number of steps [57]. Each added operation increases the likelihood of error, in keeping with what is generally known about serial processes [59].

Follette [9] extended these ideas by proposing that the quality of a particular perceptual feature to which the operation is applied should also matter. Thus, according to her model, two factors affected graph reading performance:

1) The number of operations necessary given a particular task-graph combination; and
2) The quality or effectiveness of the perceptual features used as input for the operations [60,61].

Cleveland’s ranking [62] (see Figure 4-5) was used to rank the effectiveness of perceptual features. Higher-ranked perceptual features should lead to better performance. Mental operations act on the highest-ranking perceptual feature available within a graph. The set of mental operations proposed by Follette included basic arithmetic operations such as summation and ratio estimation.

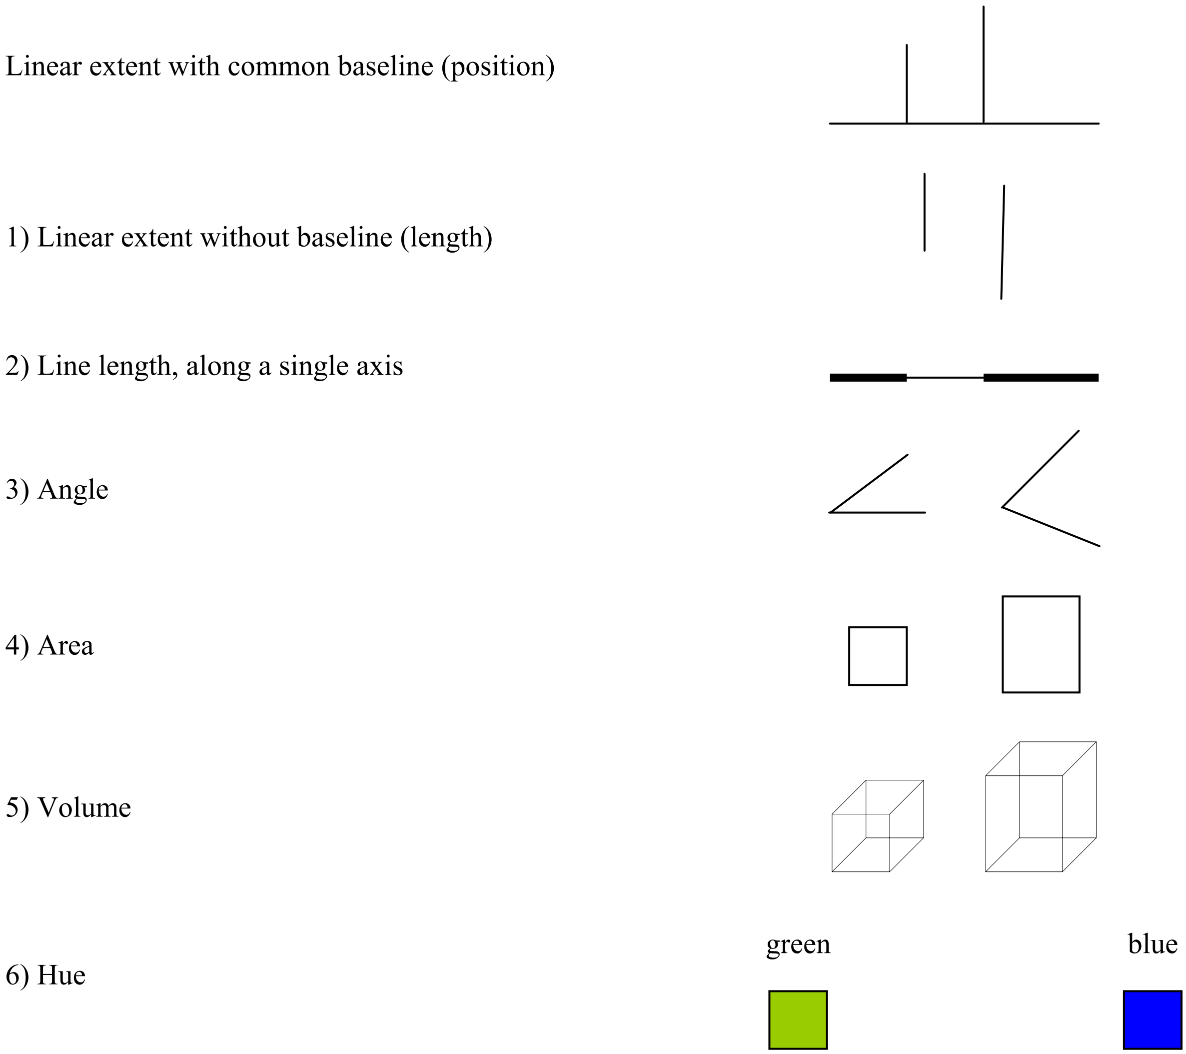

Figure 4-5: Cleveland's [61,62] Ranking of Elementary Graphical–Perception Tasks.

To illustrate the mental operations framework, reconsider the Follette and Hollands [55] results. Table 4-1 (modified from [8,62]) shows the proposed operations for each of the four conditions in their experiment. For example, part-to-part performance with graphical operations was better with bar graphs than stacked bar graphs since the one graphical operation (ratio estimation) had a more highly ranked perceptual feature (position) than did the stacked bar graph (length). Part-to-whole performance was shown to be better with stacked bars than bars since fewer operations were necessary to complete the task (no need to perform the summation operation).

Table 4-1: Follette and Hollands' [55] Results Interpreted in Terms of Mental Operations and Perceptual Features

	PART-TO-PART (A/B)	PART-TO-WHOLE [A/(A+B)]
BAR GRAPH	1) Ratio Estimation (Position); (A/B)	1) Summation (Position); 2) Ratio Estimation (Position); [A/(A/B)]
STACKED BAR GRAPH	1) Ratio Estimation (Length); (A/B)	1) Ratio Estimation (Length); [A/(A/B)]

4.4.2 Object-Based Advantages

Follette [8] examined participants' accuracy when making the two kinds of proportion judgments with several data values. With multiple values it is possible for bars involved in the judgment to be either adjacent or separated by other bars (non-adjacent). The results obtained in Follette's non-adjacent condition are of greater real-world relevance than the adjacent situation, because with more values it becomes less likely that values of interest will be adjacent. Follette found that bar graphs were less accurate than stacked bars for both tasks (although more noticeably in the part-to-part task) when the bars were not adjacent. This is surprising when considering Cleveland's hierarchy [61,62], since one would expect the bar condition to be more accurate than the stacked bar, given that all graphical elements of the bar graph are aligned on one common axis.

St-Cyr and Hollands [63] argued that Follette's [8] results may be explained by object-based theories of attention [64] which generally predict that judgments of object elements are more accurate when the elements belong to a single object rather than different ones. Similarly, in Follette's experiment, the different segments of a stacked bar are part of a single object, producing an object-based advantage relative to separate bars.

To test this hypothesis, St-Cyr and Hollands [63] created a staggered stacked bar arrangement by jittering the bars forming the stack right or left at random so that no pair of bars was contiguous. The vertical alignment of each graphical element was preserved. In their experiment, participants made part-to-part proportion judgments on non-adjacent graphical elements using bars, stacked bars, and staggered stacked bars. They found that judgments of proportions shown in staggered stacked bars were less accurate than judgments of stacked bars.

In addition, bar graphs were less accurate than stacked bar graphs, replicating Follette [8]. An analysis of the effects of distance between elements suggests that the stacked bar advantage was not due to greater distance between bars in the bar graph condition. Thus, a common baseline did not improve participants' judgments, despite the predictions of Cleveland ranking [61,62]. The results also showed that staggered stacked bars were more accurate then bars, an unexpected result. This might be explained by the fact that staggered stacked bar graphs preserve the vertical arrangement of stacked bars, although graphical elements are not part of a single object.

In the St-Cyr and Hollands [63] experiment, participants judged non-adjacent graphical elements. This is a more common situation than adjacent elements given that most graphs depict multiple elements, and the probability of adjacency is reduced with each added data point. Based on our results, this study proposes the following revision of Cleveland's ranking of perceptual tasks [61,62] for graphical perception:

1) Vertical arrangement of graphical elements in a single object;
2) Vertical arrangement of graphical elements; and
3) Graphical elements with common baseline.

Object-based advantages appear to be an important element of graphical perception, although further investigation is necessary.

4.4.3 Conclusions

Follette [8] proposed that two factors affect quantitative judgments with graphs:

1) The number of operations necessary; and
2) The effectiveness of the perceptual features used as input for the operations.

This framework was used to account for task-dependent results obtained by Follette and Hollands [55]. St-Cyr and Hollands [63] extended the framework to include object-based effects by suggesting a revised ordering to Cleveland's hierarchy of graphical elements. Although much empirical work remains to test the model's assumptions and predictions, the model offers an intriguing simple account for how graphically depicted information is processed.

4.5 THEME 4: VISUAL ATTENTION AND VISUAL SPAN

Even when a tactical display accurately depicts all relevant data, the human observer may not attend to all displayed elements. The effectiveness of Symbicon symbology to provide relevant tactical information "at a glance" was demonstrated using a change-blindness paradigm. Using an eyetracker, a gaze-contingent display can be constructed that depicts only that region of the display upon which the observer fixates. By varying the size of that region in small amounts using a staircase procedure, the result indicate how much display information can be attended "at a glance" by a human observer.

4.5.1 Detecting Change in Tactical Displays

Keillor, Thompson, Smallman, and Cowen [65] were interested in determining the degree to which different types of tactical symbology could be selectively monitored for change. The operator of a naval tactical display has multiple tasks, but the primary task is to maintain awareness and understanding of the tactical picture by monitoring changes. During such monitoring tasks, operators are susceptible to *change blindness* [66], which refers to the phenomenon that humans have difficulty in perceiving major changes in objects from one scene to another. Normally, low-level motion transients help observers direct attention toward areas of the display that rapidly change. However, if an observer blinks, glances away from the screen, or orients focal attention toward another portion of the display, the transients fail to produce a shift of attention to the changing location, and the observer is left with no awareness that a change has occurred [67]. A flicker paradigm [68] is typically employed to investigate change blindness. This methodology uses the brief flicker of a blank screen to simulate eye blinks or diversion of attention away from the changing area of the display (Figure 4-6), thereby simulating the attentional shifts that occur during naturalistic monitoring tasks.

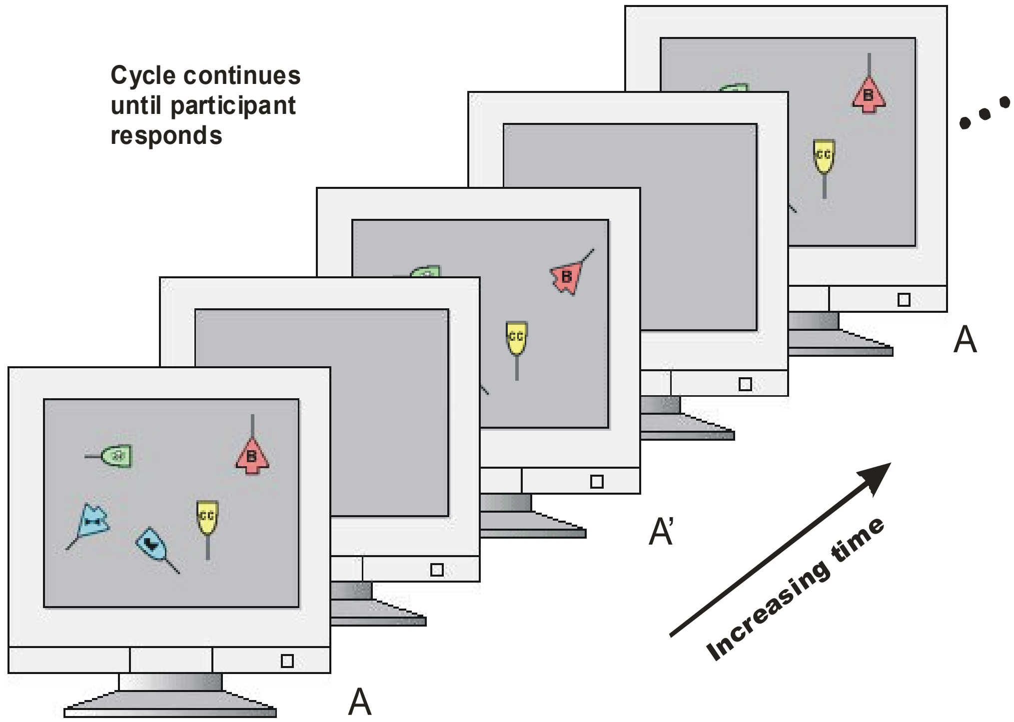

Figure 4-6: An Illustration of the Change Blindness Paradigm Used in [65].
Note the change in heading for aircraft B from screens A to A′.

Figure 4-7 illustrates a set of symbols typically used on tactical displays (MIL-STD-2525B) along with a set of symbols called Symbicons recently developed by the United States Navy [69]. Symbicons were designed to incorporate the platform (centre letter/icon) and threat (colour and frame) information of conventional military symbols with the platform classification (frame) and heading (leader line) information of realistic icons in order to promote the rapid appreciation of all depicted track attributes. Thus the main differences between the two sets are the frame shape and the depiction of heading. For MIL-STD-2525B, heading is indicated by the orientation of a leader line, but in the case of the Symbicons, the entire symbol rotates along with the leader line to face the compass direction.

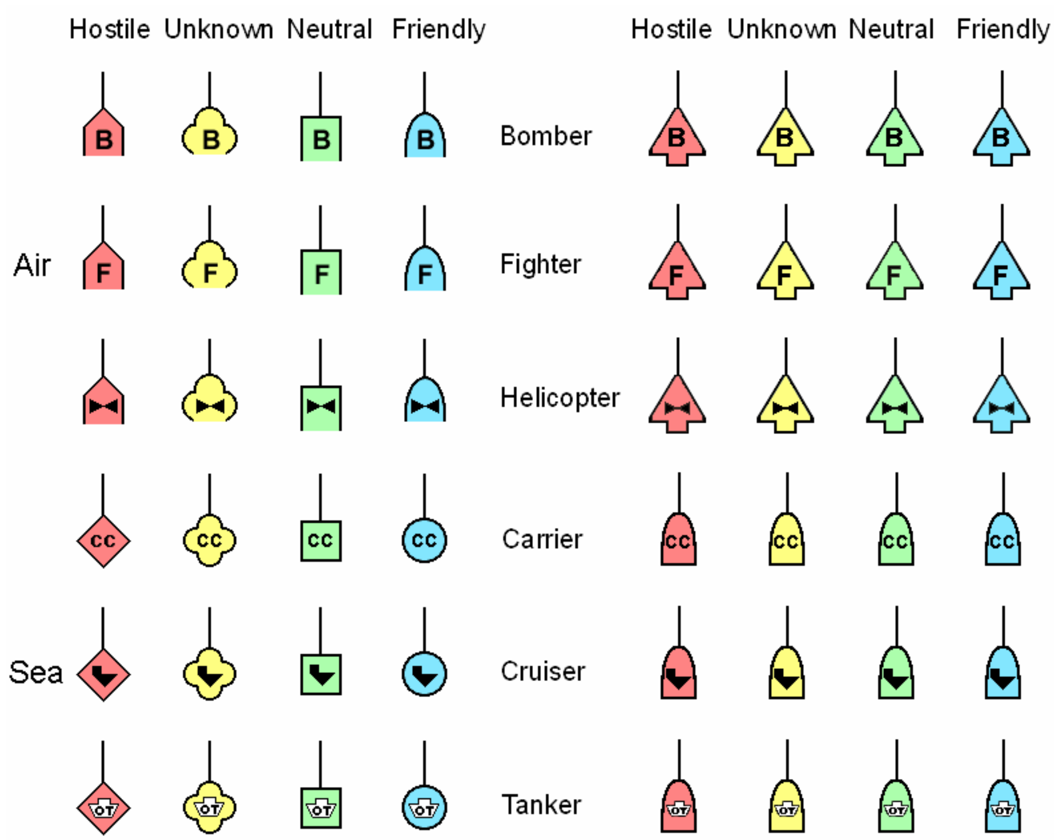

Figure 4-7: Specific Examples of MIL-STD-2525B Symbology (left) and Symbicons (right), for a Variety of Platforms and Threats.

Given the potentially disastrous consequences of change blindness, it would be beneficial if a symbology could be developed that shows greater resistance to change blindness by permitting larger numbers of contacts to be attended at once. The features of Symbicons described above may provide such resistance. Using the DRDC Toronto CVT, Keillor et al. [65] examined how quickly operators detected changes in a simulated tactical display as a function of symbology (MIL-STD-2525B vs. Symbicons).

A change detection task permits an evaluation of the extent to which operators selectively monitor subsets of relevant contacts within a display. Keillor et al. [65] examined the ability of their observers to monitor "sea" and "air" contacts for heading changes with the two symbologies. This task involved selectively attending to the frame shape (platform classification) of the symbols. For each trial, a flicker sequence comprising the original image, a grey screen, the modified image, and a grey screen was presented (Figure 4-6). The sequence continued until the participant responded "change" or "no change" by pressing the right or left button respectively.

Keillor et al. [65] found that change detection was better for Symbicons relative to MIL-STD-2525B, and that this effect was greater when participants were told in advance what type of contact (platform category) could change heading. The advantage for the Symbicons is consistent with groups of similar contacts being simultaneously attended and monitored for change.

4.5.2 Visual Span

As displays become larger or more complex, it becomes important to develop metrics to determine how much of a given display an operator can take in "at a glance". A measure of the spatial extent of visual processing is useful both in terms of optimizing display design, and developing a model for how an operator accomplishes a task when interacting with the display.

The most widely accepted way to measure visual span is by using a gaze-contingent moving window technique. By using eye position information to drive the display it is possible to obscure all information in a given display with the exception of that which is centred on the participant's fixation. By manipulating the size of this moving-window it is possible to determine how much of the display is being used, such that making more of the display visible is of no benefit. Recently, a more efficient psychophysical approach has been adopted by some researchers, in which an iterative algorithm varies the size of the window over successive trials to determine the visual span [70].

Keillor, Desai, Hollands, and Reingold [71] used a variation of the gaze-contingent algorithm [70] to examine the visual span for tactical displays as a function of the task and type of symbology used. Specifically, MIL-STD-2525B and Symbicon symbologies were evaluated. In order to determine whether visual span is affected by the task, participants were required to search for targets within an array based on threat affiliation (coded by colour), heading (coded by colour and orientation), or by the conjunction of threat affiliation and heading (coded by a combination of colour and heading).

Keillor et al. [71] found that the amount of a display processed "at a glance" is affected by the type of search task, even given identical displays. Visual span was smaller for the more complex conjunction search task. The result demonstrates that it is not the display complexity itself that underlies the span of visual processing in a search task, but instead the cognitive demands of the task. Visual span is therefore best thought of as a measure influenced by strategic and cognitive factors as well as a task's perceptual demands, making it an excellent indicator of the ease with which information may be extracted and processed from a given type of display.

4.5.3 Visual Attention and Visual Span Conclusions

In sum, the visual attention work conducted at CVT indicates that the Symbicon symbology which encodes platform classification and heading more iconically rather than symbolically better enabled human observers to detect change. Using an eyetracker, a gaze-contingent display was constructed that depicts only that region of the display upon which the observer fixates. The visual span for an observer conducting visual search was smaller for more complex conjunction searches.

4.6 CONCLUSIONS

IST-021/RTG-007 considered visualisation technologies, including display devices and techniques in relation to how they help humans to perform their tasks effectively. The RTG emphasized the human use of the computational subsystem in ensuring that the right information is available in the form and at the time needed. The experiments conducted on the DRDC Toronto CVT were essentially attempts to empirically assess those factors that make information available in the right form given task constraints and data requirements.

Specifically, a commander's tasks may require both egocentric and exocentric views on a battlespace. We examined the effectiveness of two visual momentum methods to assist the user switch between these diverse views of geospatial data. The first method (a compromise method) involved the use of a tethered display and showed that it was an effective method for depicting one's egocentric position with respect to other key locations in the terrain, which requires a certain exocentricity not available in the completely immersed view. The second method (a transition method) involved depicting smooth rotation from a 2D to a 3D terrain view, and was shown to be more effective than a discrete shift between 2D and 3D views.

We examined the use of reference points to improve the accuracy of quantitative judgments in information displays. First, having more reference points available in the display can improve performance. Second, the particular response method used, which would probably be driven by task demands, turned out to have a clear effect. Third, the use of linear perspective to depict bars at different distances from an observer introduces an source of error in judgment error and should be avoided to the extent possible.

We described work that assessed the utility of a model of mental operations in graphical perception. This model was used to account for a set of data and extended to include object-based effects. From a design perspective, the intent is to shorten sequences of perceptual operations conducted on graphical elements of the highest quality according to Cleveland's [61,62] ranking.

Finally, we described work assessing how a human observer attends to information on a tactical display. This was done in two ways: first by using the change blindness paradigm to assess what display factors affected how quickly an observer could detect a change in a tactical display; second, by using a gaze-contingent display to assess how the operator's visual span is affected by changes in the complexity of an operator's task. Such methods should ultimately allow a determination of those factors leading to more of the display surface being processed by a human observer.

In future work we plan to implement many of these display concepts into prototype systems for navy, army and joint command and control systems.

4.7 REFERENCES

[1] Chief of Maritime Staff (Canada) (2000). The Naval C4I Way Ahead: A Blueprint for 2010. Ottawa: Department of National Defence.

[2] RAdm Buck (1999). Plenary Session, *TTCP Symposium on Co-ordinated Battlespace Management.*

[3] Department of National Defence (Canada) (1999). *Defence Strategy 2020: Shaping the future of the Canadian Forces.* Ottawa: Author.

[4] Naval Command and Control Way Ahead (Chief of Maritime Staff, January 2000).

[5] Gillan, D.J., Wickens, C.D., Hollands, J.G. and Carswell, C.M. (1998). Guidelines for presenting quantitative data in HFES Publications. *Human Factors, 40,* 28-41.

[6] Lewandowsky, S. and Behrens, J.T. (1999). Statistical graphs and maps. In F. Durso et al. (Eds.), *Handbook of Applied Cognition*. W. Sussex, UK: Wiley.

[7] Woods, D.D. (1984). Visual momentum: a concept to improve the cognitive coupling of person and computer. *International Journal of Man-Machine Studies, 21,* 229-244.

[8] Hollands, J.G. and Dyre, B.P. (2000). Bias in proportion judgments: The cyclical power model. *Psychological Review, 107,* 500-524.

[9] Follette, L. (2000). *Comparing the effectiveness of bar and stacked bar graphs for judgments of proportion.* Unpublished Master's thesis, University of Idaho, Moscow, Idaho.

[10] Beno, E., Dubé, R., Duval, A., Holmes, W., Pelletier, C., Thériault, G. and Thibault, G. (2004). Concept of operations for collaborative working. Draft Contract Report (v.2.2), Valcartier, Quebec: DRDC Valcartier System of Systems Section.

[11] Hollands, J.G., Lamb, M. and Keillor, J.M. (2004, September). *Modeling 3D frames of reference for the common operational picture*. Workshop presentation, NATO IST-043/RWS-006, Visualisation and the Common Operational Picture [VizCOP], Canadian Forces College, Toronto, Ontario.

[12] Amorim, M.-A. (2003). "What is my avatar seeing?": The co-ordination of "out-of-body" and "embodied" perspectives for scene recognition across views. *Visual Cognition, 10,* 157-199.

[13] Milgram, P. and Colquhoun, H., Jr. (1999). A taxonomy of real and virtual world display integration. In Y.O.A.H. Tamura (Ed.), *Mixed reality – Merging real and virtual worlds* (pp. 1-16). Tokyo and Berlin: Ohmsha & Springer-Verlag.

[14] Wickens, C.D. and Hollands, J.G. (2000). *Engineering psychology and human performance* (3rd Ed.). Upper Saddle River, NJ: Prentice-Hall.

[15] Shepard, R. and Metzler, J. (1971). Mental rotation of three-dimensional objects. *Science, 171,* 701-703.

[16] Trafton, J.G., Trickett, S.B. and Mintz, F.E. (2001). Overlaying images: Spatial transformations of complex visualizations. Model Based Reasoning.

[17] Trafton, J.G., Trickett, S.B. and Mintz, F.E. (in press). Connecting internal and external representations: Spatial transformations of scientific visualizations. *Foundations of Science.*

[18] St. John, M., Cowen, M.B., Smallman, H.S. and Oonk, H.M. (2001). The use of 2D and 3D displays for shape-understanding versus relative-position tasks. *Human Factors, 43,* 79-98.

[19] Colquhoun, H. and Milgram, P. (2000). Dynamic Tethering for Enhanced Remote Control and Navigation. In Proceedings of the Human Factors and Ergonomics Society 44th Annual Meeting (pp. 146-149). Santa Monica, CA: Human Factors and Ergonomics Society.

[20] Wickens, C.D. and Prevett, T.T. (1995). Exploring the dimensions of egocentricity in aircraft navigation displays. *Journal of Experimental Psychology: Applied, 1,* 110-135.

[21] Wang, W. and Milgram, P. (2001). Dynamic viewpoint tethering for navigation in large-scale virtual environments. In *Proceedings of the Human Factors and Ergonomics Society – 45th Annual Meeting* (pp. 1862-1866). Santa Monica, CA: Human Factors and Ergonomics Society.

[22] Yu, P.C. and Andre, A.D., (1999). A simulator evaluation of four perspective scene viewpoints. In *Proceedings of the Silicon Valley Ergonomics Conference and Exposition-ErgoCon 99* (pp. 233-242).

[23] Lamb, M. and Hollands, J.G. (2005) *Viewpoint tethering in complex terrain navigation and awareness.* Submitted to Human Factors and Ergonomics Society 49th Annual Meeting.

[24] Hollands, J.G. and Ivanovic, N. (2004). Task switching with 2D and 3D displays of geographic terrain: The role of visual momentum. In *RTO Meeting Proceedings 50, Massive Military Data Fusion and Visualisation: Users Talk with Developers* [Papers presented at the RTO Information Systems Technology Panel (IST) Workshop held at the Norwegian Defence Logistics and Management College, Halden, Norway 10-13 September 2002, RTO-MP-105, AC/323(IST-036)TP/20], Neuilly-sur-Seine, France: NATO Research and Technology Organisation.

[25] Hollands, J.G., Ivanovic, N. and Enomoto, Y. (2003). Visual momentum and task switching with 2D & 3D displays of geographic terrain. In *Proceedings of the Human Factors and Ergonomics Society – 47th Annual Meeting* (pp. 1620-1624). Santa Monica, CA: Human Factors and Ergonomics Society.

[26] Hollands, J.G., Pavlovic, N.J., Enomoto, Y. and Jiang, H. (2004). Advantage for visual momentum not based on preview. In *Proceedings of the Human Factors and Ergonomics Society – 48th Annual Meeting* (pp. 1803-1807). Santa Monica, CA: Human Factors and Ergonomics Society.

[27] Hochberg, J. and Brooks, V. (1978). Film cutting and visual momentum. In J.W. Senders, D.F. Fisher and R.A. Monty (Eds.), *Eye movements and the higher psychological functions* (pp. 293-313). Hillsdale, NJ: Erlbaum.

[28] Andre, A.D., Wickens, C.D., Moorman, L. and Boschelli, M.M. (1991). Display formatting techniques for improving situation awareness in the aircraft cockpit. *International Journal of Aviation Psychology, 1,* 205-218.

[29] Aretz, A.J. (1991). The design of electronic map displays. *Human Factors, 33,* 85-101.

[30] Bederson, B.B., Hollan, J.D., Stewart, J., Rogers, D., Vick, D., Ring, L., Grose, E. and Forsythe, C. (1998). A zooming web browser. In C. Forsythe, E. Grose and J. Ratner (Eds.), *Human factors and web development* (pp. 255-266). Mahwah, NJ: Erlbaum.

[31] Roth, S.F., Chuah, M.C., Kerpedjiev, S., Kolojejchick, J. and Lucas, P. (1997). Toward an information visualisation workspace: Combining multiple means of expression. *Human-Computer Interaction, 12,* 131-186.

[32] Neale, D.C. (1996). Spatial perception in desktop virtual environments. In *Proceedings of the Human Factors and Ergonomics Society – 40th Annual Meeting* (pp. 1117-1121). Santa Monica, CA: Human Factors and Ergonomics Society.

[33] Neale, D.C. (1997). Factors influencing spatial awareness and orientation in desktop virtual environments. In *Proceedings of the Human Factors and Ergonomics Society – 41st Annual Meeting* (pp. 1278-1282). Santa Monica, CA: Human Factors and Ergonomics Society.

[34] Olmos, O., Liang, C.C. and Wickens, C.D. (1997). Electronic map evaluation in simulated visual meteorological conditions. *International Journal of Aviation Psychology, 7,* 37-66.

[35] Olmos, O., Wickens, C.D. and Chudy, A. (2000) Tactical displays for combat awareness: An examination of dimensionality and frame of reference concepts and the application of cognitive engineering. *International Journal of Aviation Psychology 10,* 247-271.

[36] St. John, M., Smallman, H.S. and Cowen, M.B. (2002) *Evaluation of schemes for combining 2-D and 3-D views of terrain for a tactical routing task*. Pacific Science & Engineering Group Technical Report.

[37] Barnes, M.J. (2003). *The human dimension of battlespace visualisation: Research and design issues.* Army Research Laboratory Technical Report (ARL-TR-2885). Aberdeen Proving Ground, MD: Army Research Laboratory.

[38] MultiGen-Paradigm (2001a). Creator (v. 2.5). [Computer software]. San Jose, CA: Author.

[39] MultiGen-Paradigm (2001b). Vega (v. 2.5). [Computer software]. San Jose, CA: Author.

[40] Guth, P. (2001). MICRODEM (v. 5.1). [Computer software]. Annapolis, MD: Oceanography Department, U.S. Naval Academy.

[41] Erlick, D.E. (1964) Absolute judgments of discrete quantities randomly distributed over time. *Journal of Experimental Psychology, 67,* 475-482.

[42] Varey, C.A., Mellers, B.A. and Birnbaum, M.H. (1990). Judgments of proportion. *Journal of Experimental Psychology: Human Perception and Performance, 16,* 613-625.

[43] Spence, I. and Krizel, P. (1994). Children's perception of proportion in graphs. *Child Development, 65,* 1193-1213.

[44] Huttenlocher, J., Hedges, L.V. and Duncan, S. (1991). Categories and particulars: Prototype effects in estimating spatial location. *Psychological Review, 98,* 352-376.

[45] Taylor, M.M. (1961). Effect of anchoring and distance perception on the reproduction of forms. *Perceptual and Motor Skills, 12,* 203-239.

[46] Stevens, S.S. (1975). *Psychophysics.* New York: Wiley.

[47] Spence, I. (1990). Visual psychophysics of simple graphical elements. *Journal of Experimental Psychology: Human Perception and Performance, 16,* 683-692.

[48] Hollands, J.G. and Dyre, B.P. (1997). Bias in proportion judgments with pie charts: The cyclical power model. In *Proceedings of the Human Factors and Ergonomics Society – 41st Annual Meeting* (pp. 1357-1361). Santa Monica, CA: Human Factors and Ergonomics Society.

[49] Rayson, J.K. (1999). Aggregate towers: Scale sensitive visualization and decluttering of geospatial Data. *Proceedings IEEE Symposium on Information Visualization* (pp. 92-99). San Francisco: IEEE Computer Society.

[50] Ely, J.E., Thomson, R.M. and Orlansky, J. (1963). Design of controls. In C.T. Morgan, J.S. Cook, A. Chapanis and M.W. Lund (Eds.), *Human engineering guide to equipment design* (pp. 247-280). New York: McGraw-Hill.

[51] Fitts, P.M. and Seeger, C.M. (1953). S-R compatibility: Spatial characteristics of stimulus and response codes. *Journal of Experimental Psychology, 46,* 199-210.

[52] Morton, A. and Hollands, J.G. (2003). Effects of response method on error and cyclical bias in proportion judgments. In *Proceedings of the Human Factors and Ergonomics Society – 47th Annual Meeting* (pp. 1654-1658). Santa Monica, CA: Human Factors and Ergonomics Society.

[53] Hollands, J.G., Parker, H.A. and Morton, A. (2002). Judgments of 3D bars in depth. In *Proceedings of the Human Factors and Ergonomics Society – 46th Annual Meeting* (pp. 1565-1569). Santa Monica, CA: Human Factors and Ergonomics Society.

[54] Hollands, J.G. and Spence, I. (1992). Judgments of change and proportion in graphical perception. *Human Factors, 35,* 313-334.

[55] Follette, L.J. and Hollands, J.G. (1998, October). Bars versus stacked bars: Effectiveness is task dependent. Poster presented at the Annual Meeting of the Human Factors and Ergonomics Society, Chicago, Illinois. In *Proceedings of the Human Factors and Ergonomics Society – 42nd Annual Meeting* (p. 1652). Santa Monica, CA: Human Factors and Ergonomics Society.

[56] Hollands, J.G. and Spence, I. (1994). Mental operations in graphical perception. In *Proceedings of the 12th Triennial Congress of the International Ergonomics Association* (Vol. 6, pp. 235-237). Mississauga, ON, Canada: Human Factors Association of Canada.

[57] Gillan, D.J. and Lewis, R. (1994). A componential model of human interaction with graphs: 1. Linear regression modeling. *Human Factors, 36,* 419-440.

[58] Lohse, G.L. (1993). A cognitive model for understanding graphical perception. *Human-Computer Interaction, 8,* 353-388.

[59] Swain, A.D. (1990). Human reliability analysis: Need, status, trends and limitations. *Reliability Engineering and System Safety, 29,* 301-313.

[60] Casner, S.M. (1991). A task-analytic approach to the automated design of graphic presentations. *ACM Transactions on Graphics, 10,* 111-151.

[61] Cleveland, W.S. and McGill, R. (1984). Graphical perception: Theory, experimentation, and application to the development of graphical methods. *Journal of the American Statistical Association, 79,* 531-554.

[62] Cleveland, W.S. (1985). *The elements of graphing data.* Pacific Grove, CA: Wadsworth.

[63] St-Cyr, O. and Hollands, J.G. (2003). Judgments of proportion with graphs: Object-based advantages. In *Proceedings of the Human Factors and Ergonomics Society – 47th Annual Meeting* (pp. 1659-1662). Santa Monica, CA: Human Factors and Ergonomics Society.

[64] Duncan, J. (1984). Selective attention and the organization of visual information. *Journal of Experimental Psychology: General, 113,* 501-517.

[65] Keillor, J., Thompson, L.K., Smallman, H.S. and Cowen, M.B. (2003) Detection of heading changes in tactical displays: A comparison of two symbol sets. In *Proceedings of the Human Factors and Ergonomics Society – 47th Annual Meeting* (pp. 1610-1614). Santa Monica, CA: Human Factors and Ergonomics Society.

[66] DiVita, J.C. and Nugent, W. (2000). Managing critical events in a multimodal watchstation. *Proceedings of the IEA 2000/HFES 2000 Congress. 6,* 469-472. Santa Monica, CA: Human Factors and Ergonomics Society.

[67] Rensink, R.A., O'Regan, J.K. and Clark, J.J. (1997). To see or not to see: The need for attention to perceive changes in scenes. *Psychological Science, 8,* 368-373.

[68] Rensink, R.A. (2000). Visual search for change: A probe into the nature of attentional processing. *Visual Cognition, 7,* 345-376.

[69] Smallman, H.S., St. John, M., Oonk, H.M. and Cowen, M.B. (2001). 'SYMBICONS': A hybrid symbology that combines the best elements of SYMBols and ICONS. In *Proceedings of the 45th Annual Meeting of the Human Factors and Ergonomics Society* (pp. 110-114). Human Factors and Ergonomics Society, Santa Monica, CA.

[70] Reingold, E.M., Charness, N., Pomplun, M. and Stampe, D.M. (2001). Visual span in expert chess players: Evidence from eye movements. *Psychological Science, 12,* 48-55.

[71] Keillor, J.M., Desai, I., Hollands, J.G. and Reingold, E.M. (2004). The span of visual processing in tactical displays. In *Proceedings of the Human Factors and Ergonomics Society – 48th Annual Meeting* (pp. 1798-1802). Santa Monica, CA: Human Factors and Ergonomics Society.

Chapter 5 – WORKSHOPS

5.1 INTRODUCTION

Two NATO workshops and two NX (Network of Experts) workshops were held under the auspices of IST-021/RTG-007. The NATO workshops were held in Halden, Norway (2002) and in Toronto, Canada (2004). An earlier NATO workshop in Canada had been sponsored in 2000 by the predecessor group, IST-013/RTG-002. NX workshops were held in 2001 in Aalborg, Denmark, and in 2003 at the Penn State University, State College, Pennsylvania, USA.

Both the NATO and the NX Workshops were intended primarily as venues at which military personnel could work with researchers and developers on visualisation issues in a semiformal but intense working environment. It proved difficult, however, for the organizers to attract as many serving officers as had been hoped. Nevertheless, because of the expertise and engaged interest of those officers who attended, much of the anticipated cross-fertilization did take place.

The format of each workshop emphasized interaction and discussion. Before the Halden Workshop in 2002, each workshop consisted of equal time devoted to "provocations" and discussions. A "provocation" was a formal presentation intended to expose some topic about which reasonable questions require answers, whether it be a topic derived from operational requirements, a research topic, or a question of system development. At Halden and thereafter, each workshop included not only the plenary provocations and discussion periods, but also devoted about half its time to "breakout" syndicates. A syndicate consisted of five to seven participants who were given a topic or issue to analyze early in the workshop, on which they presented their results at the end of the workshop. Several of these syndicate reports produced important insights, or led to recommendations for future studies.

The unofficial records of these workshops are available on the Web at http://www.visn-x.net for the NX workshops and http://www.vistg.net for the NATO workshops. The official publications of the 2000 and 2002 NATO Workshops are available through the RTO Web site at http://www.rto.nato.int with publication numbers RTO-MP-050 and RTO-MP-105 respectively.

The summary conclusions of the 2000 Workshop (published as RTO-MP-050) were:

- Evaluation is an important area of research;
- The importance of the interactive aspect rather than simply the presentation was brought out clearly;
- Simple displays can be useful in dealing with complex data;
- Modular and componentware structures ease the development on the computer side;
- User involvement needs to continue throughout the development process;
- visualisation of relationships is an important unsolved problem; and
- Support for the user's innovation, initiative, intuition, and creativity (I3C) is important, especially in the face of anomalous conditions.

Subsequent workshops built on and confirmed these conclusions, though the importance of the visualisation of relationships and networks, among other issues, became more prominent in later workshops. Though much significant work was presented and done at the N-X Workshops, it is not reported here. Their Proceedings can be obtained through http://www.visn-x.net. This chapter concerns only the official NATO Workshops.

5.2 THE 2002 "HALDEN WORKSHOP" (IST-036/RWS-005)

At the 2002 Workshop, there was much discussion of the relationship between data fusion and visualisation. The full Proceedings have been published as RTO-MP-105. The Executive Summary of that workshop said:

> *Scientists in data fusion and visualisation seemed to find mutual inspiration for future cooperation at this workshop. Both address human "information overload", but by disparate approaches. Data fusion uses algorithmic data reduction, whereas visualisation deals with displaying the data in comprehensible form. The presentations by data fusion heavyweights from the USA, among others, indicated that rather than just presenting the outputs of data fusion algorithms; visualisation techniques that improve the human ability to direct sensor- and computer-resources might be required.*

Here again, the importance of interaction was made clear, this time in connection with the acquisition of the data, as well as in directing its display. The syndicates (small working groups devoted to specific topics) also made this clear, most particularly in the syndicate devoted to responses to a terrorist attack and in two syndicates that independently addressed the topic of visualising information source availability and reliability.

Several presenters at the workshop considered the interface between visualisation and data fusion. Fusion means the combined use of data from different sources in a single interpretation. This necessarily means that the higher-level interpretation is of a different character than the data from any single source. At very low levels, radar imagery may be used together with geo-registered photo-imagery or multispectral sensor imagery to produce pictures in which each pixel contains some information from all the sources, and in which the patterns that manifest themselves over regions of the picture might not be obvious in the pictures from any single source.

For data that are at heart geospatial, or like military dispositions are located on a map, it is often convenient to be able to display a very large number of pixels on a large display surface. Such a surface allows for multiple viewers, it allows viewers to focus naturally on different aspects of a situation without invoking deliberate mechanisms such as mouse movements, and it requires no special tricks to help users retain context. The large display may be on a wall-like screen, or on a table-like surface as in the "CODS" display demonstrated by the Norwegian keynote speaker.

The two modes of display have different social characteristics. The wall-like display is well suited for a briefing, in which a primary presenter shows material to a seated audience, whereas the table-like display is more suited to a planning session in which several active participants range themselves around the table and are equally able to point to elements of the display. The concepts of "up" and "down" on the wall-like display have implications for the visual accessibility of different elements of the displayed situation. Those concepts do not apply to the table-like display, which might be advantageous in some cases and disadvantageous in others.

More typically, data sources cannot be spatially registered. Some may be linguistic, while others are imagic – i.e. in the form of imagery. In this case, an intermediate or common form must be found in which both kinds of data can be represented. Several talkers addressed this issue, describing ontologies, the recovery of spatial data from linguistic forms, and the linguistic representation of spatial data. From these representations, different kinds of display can be constructed. All of this work is in highly experimental stages, but would be interesting to pursue in other forums.

5.2.1 Counter-Terrorist Special Operations

One syndicate considered visualisation for Special Operations in the context of a terrorist event. The syndicated noted:

> *The steps to responding to terrorism include;* protect, detect, react and restore*. Protect involves building defensive measures. For example in an embassy you would install bomb-proof windows and place barriers around the building. Detection includes an active intelligence gathering capability. It would also include searches such as what we are subjected to at airports. Reaction is your response to the act. We considered a physical terrorist act for which the response would be a special operations force, be they military or police. Restoration would involve business continuity and/or continuance of government.*
>
> *We focused on reaction.*

The syndicate considered several different kinds of scenario, including the elimination of a sniper holed up in a building, an explosion with release of toxic material into the atmosphere, a hijacking incident, as well as responses to other civil disasters (which might or might not be the result of terrorist activity. These cases have certain fundamental requirements in common, though the importance of each requirement differs according to the situation: Navigation, Communications, Information (including Intelligence), Sensory enhancement, and Protective Equipment.

In the "Sniper Scenario" (Figure 5-1), a sniper was operating from a normal house. To stop the sniper, the authorities had two options, to invade the house or to shoot the sniper from outside when he showed himself at a window.

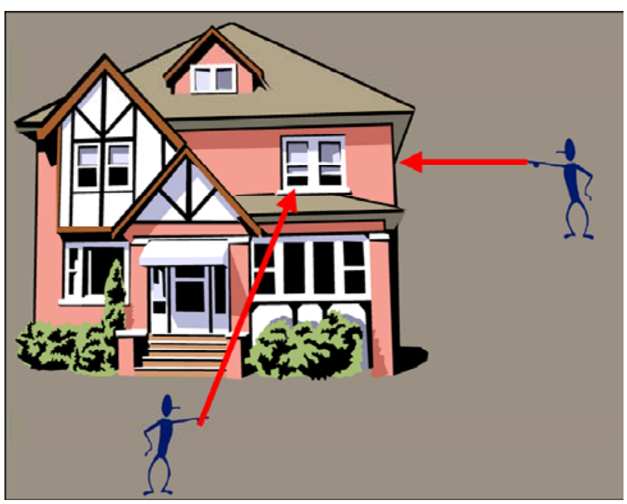

Figure 5-1: The Sniper Scenario.

Considering the list of requirements above, the authorities might need to be able to navigate not only to the house without exposing personnel unnecessarily to sniper fire, but also in the immediate neighbourhood and even inside the house itself. The operational personnel would need to be able to communicate among themselves and with their base, as well as possibly to call up other resources and information such as the construction details of the house.

This latter is one of the elements of the "Information" requirement. Another such element might be information about the social and religious background of the sniper, and of the weaponry to which the sniper had access. In the area of Sensory Enhancement, the authorities might use laser range finders and GPS to create for each officer a real-time image overlay onto a 3D plan of the house construction, in which the location of cooperating officers might be indicated. Line of sight calculations indicating areas not visible to the sniper might also be included in these composite images.

Protective equipment is included in the list of requirements. In the "sniper scenario", the need is obvious and demands no analysis.

The second scenario concerns an explosion at an industrial site that releases toxic chemicals or gases into the air. It is assumed that no live terrorists remain present, but the site must be secured, vulnerable neighbourhood residents evacuated, and clean-up procedures initiated. In this scenario, the same five requirements hold, with a greater emphasis on Protective Equipment against possible toxic agents. Communications and Intelligence also gain in importance, particularly as it concerns both the protection of local civilians and the possibility that active terrorist agents might still be present near the scene, or that other explosive or disruptive devices might be dormant, awaiting activation.

The syndicate produced a draft list of generic Human Factors issues related to visualisation in Special Operations. These included:

Presentation

- Navigation
- Symbology/maps/graphs/text
- Readability/consistency/uniformity/colour/texture

Communication

- Auditory/visual
- Voice recording/voice exchange
- Increase trust and confidence by
- Providing feedback
- Clarify source of information/reliability
- Consider absence of individual from the team
- Prevent unnecessary distraction/task interruption
- Minimize alerts

Information (including Intelligence)

- Information compression in order to avoid overload
- Reduction of vulnerability of the information network

Human Sensory Enhancement

- Facilitate perception

Equipment

- Physiological measures (Visualisation aspects for the Command Centre)
- Ergonomic aspects (anthropometrical data, visibility, reachability, colour)

Training and Preparation

- The more the HF aspects are taken into account the less training is necessary.
- Task- and user-oriented

Planning Aids

- Support hierarchical planning (rough/detailed)
- Provide feedback
- Apply heuristics

Real-Time Translations

- Provide messages/info as universal as possible
- Provide messages/info as simple as possible
- Individual translation if necessary

The full syndicate report is available in the Workshop Proceedings (RTO-MP-105). The topic was considered sufficiently important to warrant proposing an Exploratory Group to prepare an RTG.

5.2.2 Visualising Potential Information Sources

One of the dilemmas facing commanders in fluid situations is that information sources of possible immediate value may exist without the commander knowing of their existence, or of their reliability if their existence is known. These resources can include national, coalition, theater, organic, or global sensors, systems and platforms, or other data sources (open press, media, www, etc.).

Any complex scenario will afford conventional sources of relevant data, but in addition there will probably be less obvious sources that could be very useful if only their existence and quality were known. Two syndicates at the workshop considered this problem. The following section is extracted and edited from the report of one of these syndicates. It used the VisTG Reference Model as a framework for considering the issues.

5.2.2.1 Using the VisTG Reference Model in the Scenario

Several times during the workshop it was mentioned that one of the problems with coalition operations is that nations do not necessarily make pass to the coalition command all the information they have. Some of this information might be made available on request, some might be made available after being sanitized, and some might not be made available outside the national command. Furthermore, the command staff may find for particular situations that unconventional information sources are both available and useful. Hightempo decision makers require an awareness and understanding of the existence, availability, and quality of time-dependent information resources.

The question addressed by the syndicate was: *How might a visualisation system assist in the understanding of the existence, availability, and quality of information resources in a distributed NATO or coalition operational environment?* It was assumed that very many different kinds of data source might be available, both official and unofficial, and that the nature of all of them would be contained in some kind of a dataspace accessible to the user (the command staff). There would be too many possible data sources for any staff officer to be able to remember all of them reliably.

To assist the analysis and understanding of this problem, the group explored the potential situation of a downed pilot during a tactical mission tasking. The tactical operational commander needs to obtain the necessary information in order to plan and execute an extraction operation. To make the commander's problem explicit, the syndicate proposed a possible chain-of-events, some of which follows:

1) Pilot goes down.
2) Transponder on aircraft is activated.
3) Table console (TC) in command and control (C&C) flashes red dot.
4) Commander zooms TC into pilot location to view terrain/image model at fine resolution.
5) J2 and J3 initiate a data availability request for area of interest (AOI) around the pilot.
6) Data list is retrieved from metadatabase that is continuously updated by Allies. Such a list might include, along with non-conventional public and private data sources:

Source Platform	**Available Data**	**ETA**	**NIC**
UAV	High Resolution IR	5 min	US-12
Platoon	Visual	30 min	NOR-2
Satellite	1 m Multispectral	3 hours	CAN-3
Special Forces	Medical Condition	1 hour	UK-1
F-16	Sighting	-10 min	FRA-4

7) … (several stages omitted).
8) Display the "probability of detection" or "probability missing data" for any variable the commander specifies.
9) Predict and display the enemy forces' situational awareness of AOI.
10) Enable ability to gain control of, and task, mobile sensors.

In order to conduct an analysis of this problem, the group chose to follow the VisTG Reference Model of *visualis*ation system design (Figure 5-2). The design process using the VisTG Reference Model proceeds in stages. The model assumes a series of nested loops, each based around the achievement of some purpose, or goal, and the perceptions needed by the user if progress toward achieving that goal is to be assessed. At any one level, there may be several parallel loops, and any one of these may simultaneously serve more than one purpose at a higher (outer) loop level. For each of these many loops, the method identifies six basic questions, which can be summarized:

1) What is the purpose to be achieved?
2) What does the user need to perceive if the purpose is to be achieved?
3) What does the user need to be able to do to achieve the desired perception?
4) What impediments might detract from the user's ability to perceive (including lack of user training)?
5) What impediments might reduce the user's ability to act to affect the necessary perception (including lack of user training)?
6) What provisions might be available for alerting the user to portions of the dataspace potentially relevant to the purpose of the loop?

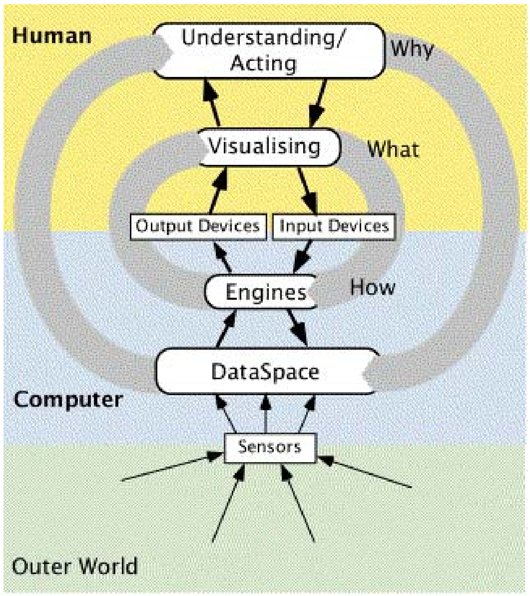

Figure 5-2: The VisTG Reference Model.

5.2.2.1.1 Outer (Level 1) Loop

The first stage of the design process, then, must be to assert the purpose that a given loop is to serve. In the problem at hand, the system to be designed is supposed to help a command staff (including the commander) to be able to understand those data sources that might reasonably be expected to have, and to be willing to provide, data relevant to the situation of immediate interest, including understanding of their probable reliability, relevance, and latency (time to deliver the information).

Translating these requirements to the use-case, relevant data sources include those that could potentially supply information about:

- The location of the downed aircraft and its pilot;
- The pilot's state of health;

- The character and socio-political environment of the people in the region;
- The location, numbers, identity, and movement of nearby blue, red and orange forces; and
- The local physical environment, such as terrain, weather, trafficability, visibility, and so forth.

The answer to Question 1 for the outer loop (i.e. Q 1.1), *What is the purpose of the system*, then, is that the system is to help the command staff to distinguish from a potentially very large number of official and unofficial data sources those that could potentially supply any of the desired kinds of information, and to prioritize among the potentially relevant ones those that would be most likely to be useful within the time frame available for the operation.

Question 1.2, *What does the user need to perceive*, is about the relationship of the sources to the necessary data. The data requirement information is paramount, the source is a means to acquire it. This means that the user needs to perceive the source as an attribute of the required data, not the reverse. The system must be able to present the data so that the user can see not only which sources can provide which specific kinds of data, but also the likelihood of getting the information, its probable reliability, and the probable delay (latency) before the information is made available.

Question 1.3, *What does the user need in order to be able to act to achieve the desired perception*, is about the user's ability to let the system know what kinds of data are desired, so that the system can search its dataspace to determine what sources might be able to provide them. The use has to be able to let the system know such things as the degree to which latency or reliability matters, just as much as to let the system know what kinds of data are required.

Question 1.4, *What impediments might detract from the user's ability to perceive what is necessary,* is about both external impediments, such as the uncertain willingness of National Intelligence Cells to release information they may be known to have, and internal impediments such as the user's lack of knowledge that certain kinds of data or sources of data might be accessible or useful. The latter component may be characterized as possible lack of training or experience.

Question 1.5, *What impediments might detract from the user's ability to act to generate the desired perception,* is about the user's ability to get out of the database the displays that would show the immediately useful sources, the usefulness of the data they could provide, and the means to acquire the information from the useful sources. It is a question about user control of the content of the presentation displays (sometimes erroneously called the "visualisations"). In the use-case, for example, the system might make it difficult for the command staff to specify that they needed to ascertain how long it would take to acquire imagery from a UAV that might already be in the vicinity or that might have to be tasked to take off on a specific mission related to the retrieval of the downed pilot. This information might affect the relevance level of other possible sources of similar data.

Question 1.6, *What provision is there for alerting the user to potentially useful regions of the dataspace,* concerns autonomous actions performed by the system under general rather than specific control of the user. The user (or system designer) may set up criteria for determining what patterns in the dataspace warrant being labelled "potentially useful", but the system does the labelling independently of immediate user control. For the use-case of a downed pilot, the system might be set up to highlight data sources related to the attributes mentioned above: location, terrain, weather, pilot's health, local socio-political environment, force dispositions, etc. For other scenarios, other constellations of data sources might be highlighted.

5.2.2.1.2 Engine (Level 2) Loops

The questions discussed above are all cast in terms of the purpose of the system as a whole, and as a group, they specify the requirements and illustrate some possible pitfalls to be avoided in the design. The design process builds on the answers to the Level 1 questions, using them as partial specifications for questions that define Level 2 loops. The same series of questions apply in the analysis of each loop.

5.2.2.2 Suggestions for Using VITA in the Scenario

Consideration of the VisTG Reference Model led the syndicate consider whether the VITA engine might be usefully developed for situations like that of the use case scenario. The VITA engine suite (Chapter 2) was developed in the context of discovering patterns of concepts within Web pages, but is not limited to that area of application. It contains both Data Selection and Algorithm Execution engines, and Algorithm Selection is to some extent also incorporated in the form of prior user control.

In the VITA original design, one or more "Query" instances are submitted to search engines. Each Query contains a set of concepts (keywords, in the initial implementation), and the search engines return for display Web pages that contain at least a threshold number of concepts. In the 3D VITA display, the Queries are displayed as nodes in one plane, the concepts as nodes in a second plane, and the Web pages that pass the threshold test as nodes in a third plane. Concept nodes are shown as linked to Query nodes on the one side, and to Web page nodes on the other. The placement of concept and Web page nodes within their respective planes is controlled by a self-organizing process that locates similar entities near each other.

In the context of a data source discovery process, a similar display might substitute Standard Scenarios for the Queries, Data Requirements for the Concepts, and Data Sources for the Web pages. The user's eye would tend to be drawn to those data sources most relevant for the scenario at hand. If more than one scenario applied, the display would be essentially the same, except that the "Query" plane would be more heavily populated, as is the case with "standard" VITA when displaying the combined results of several independent Queries.

The "standard" VITA display allows for variable representation of concept and page nodes according to their attributes, such as whether the page has been examined. Using a VITA-like display for Data Source discovery, valuable attributes might affect either the placement or the representation of the node, or both. For example, probable latency might affect the transparency or some shape-related iconic value of the data source representation, or might affect its location above or below its "natural" representation plane. Or, if the user were given the ability to indicate to the system the relative importance of attributes such as latency and reliability, the data sources might be represented in colour variations that represent the important attributes. Many possibilities for enhancing the standard VITA presentation to take advantage of the requirements were brought out by the "six-question" analyses generated by using the VisTG Reference Model.

5.3 NX WORKSHOP 2003 "INFORMATION VISUALISATION NEEDS FOR INTELLIGENCE AND COUNTER-TERROR"

The 2003 NX Workshop was held at Penn State University, USA, 10-11 March 2003, with the general topic area: *Information Visualisation Needs for Intelligence and Counter-Terror.* As it was not an official NATO Workshop, only a brief summary will be given here.

The workshop was attended by 29 experts from seven NATO nations. Its procedures, including the use of small working groups (Syndicates), were modelled on those of the Halden Workshop (IST-036/RWS-005). The following syndicate topics were addressed:

- Information Visualisation for Counter Intelligence;
- Information and Data Source Discovery, in a coalition context; and
- Information Taxonomy for Presentation, Selection and Design.

The work of the three syndicates resulted in two reports, one that followed on the work discussed above, and one that looked forward to the 2004 Toronto Workshop (IST-043/RWS-006):

- Visualisation needs for counter-terror-directed intelligence gathering; and
- Future needs and challenges of a common operating picture for defensive information warfare.

Presentations and reports from this and other NX Workshops are available on the web at http://www.visnx.net.

5.4 THE 2004 "TORONTO WORKSHOP" (IST-043/RWS-006)

At the 2004 Workshop "Visualisation and the Common Operational Picture", 52 military users, developers, and researchers from 9 countries discussed a comprehensive list of issues and comments relating to the Common Operational Picture (COP). At the end of the workshop, some of the findings and proposals were summarized as follows:

- NATO used manually edited PowerPoint presentations for COP in support of security at the Athens Olympic games. This was the operational state of the art, in contrast to what was at the time available in research institutions and in demonstration projects. One issue is therefore how to get into operation the best of what is actually available to the user.

- Data-bandwidth
  - There is a need to have rules to regulate what takes up the bandwidth and pre-fit information for higher levels minimizing what gets through the bandwidth.
  - The most complex problem is the management of documents.
  - Database size.

- Framework for COP
  - Other countries would not trust a country that is not being up to data in the NATO plan. Countries have been kicked out. They were in the dark and unprepared. It is political. The goal is to have 60% of everyone up to par. This will go on for years, so get the framework going.

- Standards
  - A lot of nations are going to be in the same situation. We need to determine: "Do we have a common Operating Framework". The nations can not create their own COP. You need to tell nations that if they are going to be somewhere specific, they need this and this… If they do not have it they are in the dark. We need standards. We need to spend some time discussing this.
  - Rationalisation is a better word than standardisation. People always object to submitting their own stuff to someone else's standard.
  - Be careful about how a NATO wide standard is presented. There are lots of cultural differences. What could be good for one country could be really bad for another country.

  - A COP where everyone is in and participating is needed, but there will always be differences between commanders, and countries where people say "I want this" instead of the standard. Leave some room to be flexible for the group.
  - Every country can have its own database, extend the model, and individual countries can change their own model, when a good idea arises, without all of the NATO countries having to change their plans.

- Uncertainty
  - There are several aspects. How do we communicate uncertainties to commanders? There are several visual conflicts. Should it be standardized rather than nation-specific?

- Cultural differences
  - Cultures change slowly, and missions and our need for pictures change faster, so if there are cultural differences, they can be predetermined.

- Training
  - If people have to train to use something, they would not bother. Give them something that is easy to adapt to. Have it available to them, if they want to use it. We need to spend more time on the Human Factors side of it.

- Difference between spatial and non-spatial background (which ties into the HAT-report (RTO-TR-030) distinction between *located* and *labelled* data).
- Further research
  - The need to document that one type of presentation is better for visualisation than another, after research is done.
  - Spend more time on the Human Factors side of COP.
  - Report what is possible in short term.
  - Make a technology demo with real users.
  - Make prototype for full testing.

- Different levels and kinds of commonality
- Leadership in establishing COP
  - Agreement should be made on leadership when establishing a COP.

- The biggest threats to good systems are procurement and security peoples
  - Procurement takes time. By the time the system is available the needs are different.

- Stop talking and do something
  - Show something (60 – 70% solution is better than nothing).

During the workshop, many different approaches were discussed, ranging from specific instances of how to create effective presentations for specific purposes to general theoretical positions on what was meant by a "Common Operational Picture." It was generally recognized that any two people will have different background knowledge, and that this knowledge is combined with information from a presentation in creating each person's visualisation. Logically, a "Common Operational Picture" is an impossibility. Furthermore, the roles and responsibilities of each user would ordinarily be different, which implies that they would interpret different aspects of any presentation, even if they all had the same background knowledge. What is needed, then is not a "Common Picture", but a "Common Operational Environment" within which the individual (in-the-head) visualisations coordinate effectively when planning operations.

5.4.1 An Approach to a Common Operating Environment

For each player, the operational environment includes not only the scenario (real or simulated) and the displays representing that scenario, but also all the other players. In planning any coordinated action, there are several different domains of information. One of the syndicates studied this issue. As a basic framework, they used the following motto: *A "COP" system provides purpose driven views that lead to common understanding.* The following material is taken from the report of that syndicate.

Introduction

According to the syndicate, a COP System is a system that facilitates the development of a "Common Operational Picture". The notion of a COP has three components: "Common", which implies that there are at least two collaborating partners; "Operational", which implies that there is a real-time element involving action involving the partners; and "Picture", which implies that each partner has some kind of vision of the situation in which the action takes place. This section addresses the first two of these components. The "Picture" aspect involves for the most part issues that do not change between displays intended for one user and displays intended to facilitate the development of a vision common to two or more partners.

A primary problem with the notion of the "Common Operational Picture" is that for the picture to be both "Common" and "Operational", the data on which cooperating parties base their picture must be up to date. This implies communication between the databanks on which the parties base their displays. The displays normally will not be in common, unless the parties are physically together, looking at (or listening to) the same display, but they should have sufficient commonality that each party understands the other's view of a situation well enough to be able to visualise how their respective roles in any action support one another.

To create a "vision", a person integrates incoming data with memories and understandings already in the mind. The commonalities of background data known by the partners can be enhanced by communications on widely different time scales, the immediately varying data being perhaps not very extensive (e.g. a blown bridge is easily described in a few bits of data, if the parties have a detailed reference map in common, or if the bridge in question has recently been discussed). This may seem self-evident, but it forms the basis of the syndicate's proposed approach. The underlying point is that when rapid cooperative action is required, very little data need be communicated between the partners, if they share (and know they share) an appropriate common background.

In military systems, common backgrounds arise in several ways, not least through the medium of training in the doctrine of the services to which the partners belong. If they have training and "culture" in common, communication of a "Common" vision is much easier than if they belong to different services (in Joint operations) or to different nations (in Coalition operations).

Communication bandwidth is a scarce resource during a mission, for two reasons. Firstly, all communication may be intercepted, subverted, or spoofed. The less information must be transmitted, the harder it is for the enemy to take advantage of it, or indeed, to detect it. Secondly, it takes time and resources to create an informative message, and more time and resources for the recipient to interpret it. Whatever background information can be shared will reduce the information to be transmitted, making the transmission both more secure and more quickly transformed into action.

With these issues in mind, the working group developed a Venn diagram approach to the description of information sharing among partners. Much of this section discusses the implications of the elaborated Venn diagrams.

The Basic Venn Diagrams

For simplicity, the partners who require a COP for some cooperative operation are called “A” and “B”. The basic Venn diagram is shown in Figure 5-3. Subsequent diagrams build upon this as a base.

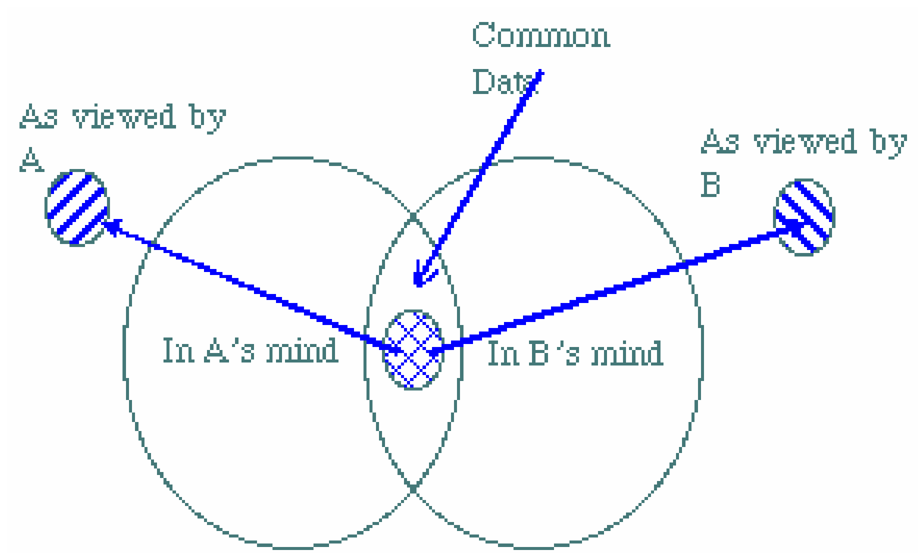

Figure 5-3: The Basic Venn Diagram. The intersection of the ellipses represents the information shared between A and B. The central textured ellipse represents the data currently being discussed between A and B, which may be displayed differently to each partner.

In Figure 5-3, the two larger ellipses represent the whole of the knowledge and memories held by either A or B. Very little of this is relevant to the operation, and even less is shared between the partners. Some aspect of the information is currently being communicated between them (the part in the textured ellipse), and that part either is currently shared, or becomes so through the discussion. If it had been intended to represent information known by A being passed to B, the textured ellipse would have been placed within A’s ellipse, outside the intersection “lozenge”, but moving into it.

In a collaborative situation, the partners will have some commonality of background, which allows for the “Common” interpretation of situations, though their differences in background information may mean that the displays needed for them to achieve a common understanding are quite different. The intersection “lozenge” indicates this common background.

Within the context of the Workshop, the COP is, at least in part, mediated through data stored electronically. The basic Venn diagram therefore must be expanded to represent the computer’s role.

It is most convenient to illustrate the computer as a third partner, which, like the humans, holds data, some of which is known to one or both of the humans, and some of which, though accessible, is not presently known to either, as suggested in Figure 5-4. Figure 5-4 shows that under discussion is some data known to the humans that is not known to the computer (as will ordinarily be the case).

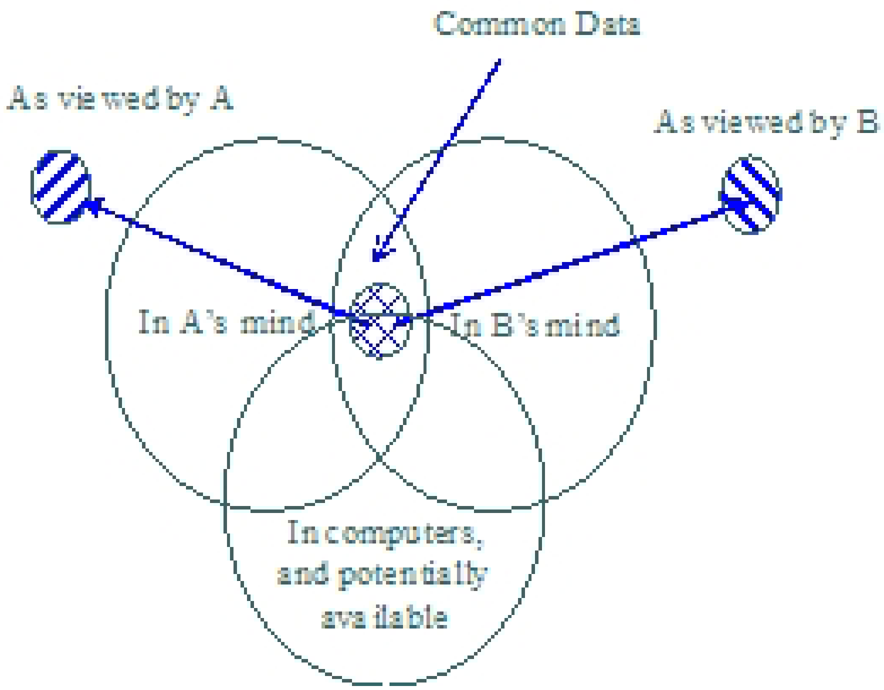

Figure 5-4: Incorporating the Computer as Intermediary and as Potential Data Source. The communication between the partners may be done both through the computer and through normal human communication media. Although the computer is shown as singular, the two partners are more likely to use different computers, linked by some data communications medium, and using shared data models.

To execute the operation successfully, the partners may need more information than is available to either, and some of it may not yet be in the electronic database, either. To represent this fact, one final elaboration of the Venn diagram is required. The grey area in Figure 5-5 represents all the information that would be needed to ensure a successful operation. Some is known to both partners and is in the database of the computer. Much of this would be factual material relating to the operation or to the doctrines within which the operation exists. Some is known to the humans, but is not in the computer. This region would include, for example, the knowledge each has of the other from having worked together on other occasions, or from having shared experiences of other kinds. It is interpersonal, and a considerable part of it may be of an intuitive nature, not easy to reduce to a form that could be computerized. Some, on the other hand, is not known to the humans or to the computer, and that is the area from which surprises emanate. As the operation evolves, it is the movement of information from that upper grey area into the space of A's or B's knowledge that will affect modifications of the operational plan.

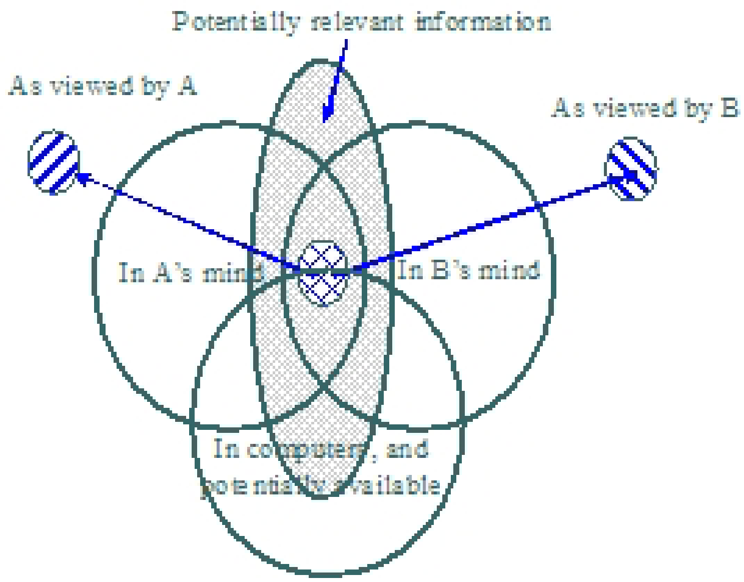

Figure 5-5: Although A and B Know much of the Potentially Relevant Data, Some of what is Known is not Shared, Some is Known to Neither, but is Accessible in the Computer Database, and Some is not Known.

Implications for the Design of a COP System

The different areas are labelled numerically in Figure 5-6, as a guide to the following discussion. In Figure 5-6, the area representing the data currently under discussion (the textured area labelled "9") has been moved to overlap parts of areas 4, 5, 7, and 8. The significance of this will be discussed below. Only areas relating to partner A have been numbered, since the discussion is completely symmetric with respect to the two partners.

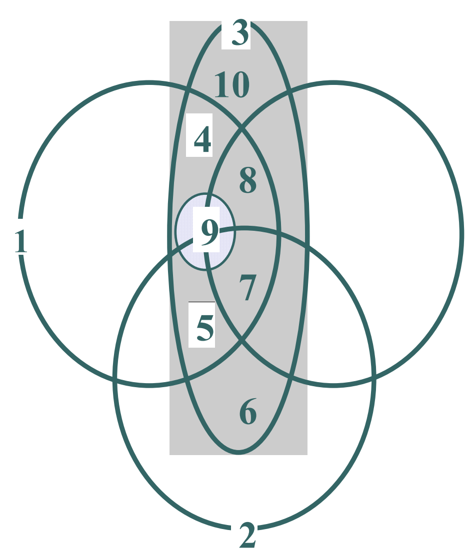

Figure 5-6: Data/Information/Knowledge Areas of Different Significance to the Developing COP.

The different areas are:

1) Everything known to A, whether relevant to the operation or not.
2) Everything potentially accessible in the computer data structures.
3) (The grey area) Everything relevant to the operation. Potentially all of area 3 could be included in the COP, but in practice it may not be necessary for B to have access to all of A's operation-relevant information.
4) Potentially relevant data known to A, but not to B or the computer.
5) Potentially relevant data known to A and in the computer's dataspace.
6) Potentially relevant data in the computer's dataspace, accessible but not known to either partner.
7) Shared information also in the computer's dataspace.
8) Shared information not in the computer's dataspace.
9) Data under discussion. Any data in area 9 that is also in area 4 or 5 would move into area 8 or 7 as a consequence of the discussion. In other words, it should become known to B, if the communication is effective.
10) Potentially relevant data not known to either partner and not in the computer's dataspace. This area represents possible danger. If either partner is aware of the possibility to seek out some of the data in this area (by tasking sensors, for example, or by querying other authorities) the data might be movable into one of the other areas – but some data will always remain in area 10, perhaps to be encountered later in the form of a welcome or unwelcome surprise.

The Venn diagram suggests that in order to create an effective COP, as much information as possible should be moved from areas 4, 5 and 6 into areas 7 and 8, though it might well be most effective for the data in area 6 (in the computer's dataspace) to be left in that area until needed, provided that the possibility of displaying that information is known to A or B. Much of this transfer between areas can, and should, be done before the operation commences and strictures on communication bandwidth become significant.

Discussion

"Provided that the possibility of displaying that information is known to A or B" indicates a significant problem with the Venn diagrams. They have no representation of whether A knows that B knows or does not know information that is shared (B does know it) or not shared (B does not know it). In other words, A and B may share knowledge of some information, which is therefore properly included in area 7 or 8, but neither may realize that the other does know it. Alternatively, A may not realize that B does not know some critical item. The COP system should have some way whereby the partners can probe each other's understanding. It incorporates a dialogue possibility beyond that suggested by the existence of the textured ellipse "9".

One of the most important types of information that should be shared is the goals, missions, or intentions of the collaborators in the operation. To take a caricature of an example, suppose A has been ordered to blow a bridge, and B has orders to cross the river spanned by the bridge. B's mission would be substantially affected if A were to blow the bridge before B crossed it! Clearly, the partners must have access to information about each other's intentions, though that information should not be allowed to clutter displays of other material germane to the operation.

The mention of goals and their disturbance immediately brings to mind the "TAGCI" (Generic Architecture and tasks for interface design) concept presented at this workshop by Heureux, Duquet and Fiset (see the Proceedings: RTO-MP-105 for a full account; a portion is extracted below). TAGCI represents a person's main goal as the root of a tree of sub goals and sub-sub goals, with which is paired another tree of potential impediments and threats that might lead to failure to achieve the corresponding goal or sub-goal. In the context of the COP, each partner's goals and subgoals present potential threats to the other's ability to achieve their goals, and thus present threats to the success of the operation for which the COP is constructed. It would therefore seem reasonable to ask whether the TAGCI approach should not be incorporated in the design of a COP system, and to form part of the structured approach to COP system design implied by the considerations in this report.

The success of the collaborative operation depends on as much as possible of the "grey area" being in areas 7 and 8. Transfer of data from areas 4, 5, and 6 into areas 7 and 8 implies communication. Clearly, incoming data during the operation, known only to A, or through electronic means to the computer, must be communicated in real time, using whatever communication channels are available at the time. What is known before the operation can be, and should be, communicated before the operation starts, and for those aspects the communication channels can be chosen freely.

For partners who have worked together on previous operations, much implicit knowledge will be shared and will be known to be shared. Subordinates will have a feel for how to interpret their commands, and the commander will have expectations as to how the subordinates will execute their missions. If the commander is new, or even more, in the case of joint or coalition operations, that kind of shared knowledge is replaced by the formalized knowledge built into the doctrines of the partners. It must therefore be supplemented by communication of the kind that builds mutual confidence.

Common background is an important part of areas 7 and 8. It is this background that allows each to know how to create and how to interpret terse communications. Common understanding of the goals further reduces the need for real-time communications bandwidth. If the commander trusts that the subordinates understand the global mission and how their goals contribute to that mission, then the commander need not be informed of their progress in real time, except insofar as it compromises or facilitates their expected contribution to the mission. These observations may be intuitively evident, but the Venn diagrams make them explicit.

Conclusion

A true "Common Operating Picture" emerges, not when the parties have the same picture, but when their different viewpoints on the situation allow them to cooperate without misunderstandings of their respective roles and functions within the operation. This happy condition may or may not come to pass when the parties see the same display, but it cannot come to pass unless they have shared access to the relevant data.

A Venn diagram captures the notion of the information known, unknown, accessible, individually known, and shared. To form the COP requires that information be transmitted from the individually known to the shared known parts of the diagram, and from the accessible to the known. For it to be effective requires that the relevant information be in the partners' vision at the operational time it is needed. To minimize the bandwidth needed at the time of the operation is an essential aspect of the construction of a COP, and that means that what information can be moved early into the shared dataspace should be moved early.

While the Venn diagram may not show anything that is not otherwise obvious, it does help the designer of a COP system to consider what information to incorporate in the COP, and when and how to move it into different parts of the shared dataspace. It also points up the importance of information shared through non-technological means, such as by shared human experiences, and indicates the value of such human-shared experience in reducing the communication bandwidths required in the press of a real-time operation.

5.4.1.1 TAGCI – A Structure of Goals and Impediments

The following is extracted from the report "*Visualisation issues in the context of Information Fusion*" by Fiset, Duquet, and Heureux, included in RTO-MP-105. TAGCI seems to complement the VisTG Reference Model, in that it provides a method of representing and handling the "impediments" that must be considered in the VisTG analysis both of the input and the output sides of each loop. It integrates these impediments across the various levels of a visualisation task.

A method called TAGCI, the French acronym for Generic Architecture and Tasks for HMI Design was developed in the process control domain [3]. Briefly, TAGCI prescribes that the monitoring and control of a complex system (and data fusion, with its large number of interconnected variable, dynamics, conflicting goals and risk, certainly qualifies as a complex system in itself) can be represented as a structure of generic tasks. Generic tasks are "basic combinations of knowledge structures and inferences strategies that are powerful for dealing with certain kinds of problems"[4]; they are used in TAGCI to derive the information requirements for an operator to accomplish his tasks. Those tasks are:

- Detection of a threat to an operating goal;
- Transition (taking the system from an operating region to another);
- Diagnosis (identifying the source of a threat to an operating goal);
- Optimization (fine tuning of the system to enhance the achievement of an operating goal within an operating region); and
- Compensation (bringing the system to a safe state in case of anomaly).

By themselves, those tasks describe how an operator (or more generically, an agent) operates a complex system or process. A useful property of generic tasks is that one can define logico-temporal constraints on their ordering. For example, detection must precede compensation or diagnosis tasks; further, detection must be constantly active to ensure that relevant goals continue to be achieved. On the other hand, transition tasks are carried out upon demand.

Since the operator understands and acts on the process through models, there is a natural relationship between the various models and the generic tasks. An analysis has shown that it is possible to map the generic tasks to specific process models. Of particular interest to data fusion is the detection task, to which is associated a type of goal hierarchy known as a goal tree success tree (GTST). The GTST provides the information required to support the detection, diagnosis and compensation tasks. Goal hierarchies have also been shown to be able to support the detection and part of the handling of expected and unexpected abnormal situations. The Figure 5-7 shows the relationship between the detection and compensation tasks and the GTST.

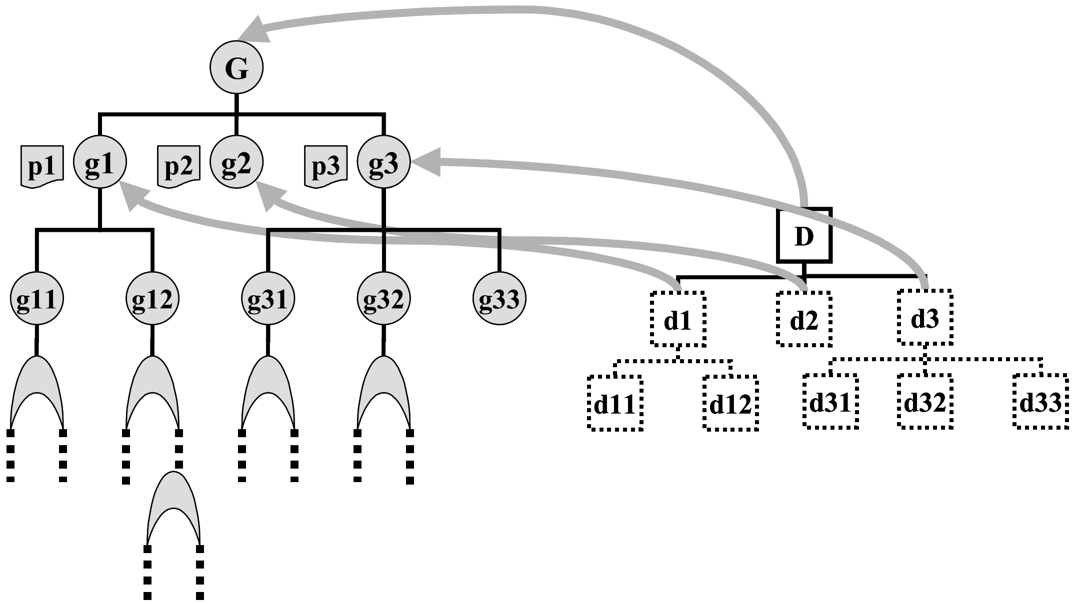

Figure 5-7: Relationships between the GTST and the Generic Tasks.

The left hand side of Figure 5-7 shows the GTST. The top part is the goal tree, which comprises the set of all domain goals that are needed to achieve a mission. The main goal (G) can be decomposed into sub-goals (respectively, g1, g2, and g3). There is an implicit AND gate between the goals and sub-goals of any given level. The bottom part of the GTST is the success tree, where success paths (connected by an OR gate) lead to the achievement of the sub-goals that in turn lead to the achievement of their parent goal(s). Compensation plans (noted p1, p2, p3 in Figure 5-7) can be found next to their associated sub-goals (which do not exclude them being present for the main goal). Compensation plans essentially hold pre-determined strategies to deal with threats to a given goal.

The right hand side of the Figure shows the detection task; this generic task is carried out by an operator (or any intelligent agent) to determine if the goal is achieved. The top level detection task (D, in the Figure) can thus be associated to the top level goal; this task can also be partitioned, if it is sufficiently complex that doing so would reduce the risk of a design error in the HMI. Those tasks are used to derive the information and control needs of the operator. It should also be obvious that with this framework, diagnosis becomes a recursive application of a detection task, down the goal tree.

Aside from the definition of a strictly technical content for the HMI, TAGCI also integrates important human characteristics. One such characteristic is the fact that operator will sometimes monitor a process through a “management by exception” strategy where he relies on various exceptions (e.g. alarms, indicators) to alert them to an active anomaly; he may also use a “management by awareness” strategy where he is actively engaged in the monitoring of key process variables in an attempt to prevent anomalies. Both types of behaviours occur naturally and may sometimes even cohabit, depending upon individual operator styles and process conditions. In terms of HMI design, recognizing those two types of behaviours, and taking into consideration typical HMI design constraints (e.g. limited display real estate) permit to identify design rules. For example, one such rule dictates that indicators supporting management by exception must always be shown, while indicators supporting management by awareness must always be available.

Further, the initial version of TAGCI includes design rules that support the selection of some simple visualisation components and guidance on how to "compose" individual displays to best support the operator.

Using TAGCI should enhance traceability of the HMI design and, more importantly, improve the performance of the "operator – data fusion" system. It also changes the way data fusion process is considered; the algorithms remain, but the operator becomes more involved into the process. However, this involvement needs to be mediated in such a way that his workload will not increase.

5.5 NETWORK REPRESENTATION

The Workshop included many inspiring and informative presentations and discussions, several of which considered issues of network representation or interaction between the user and the presentation system.

5.5.1 Types of Physical and Logical Networks

At the Halden workshop (IST-020, 2002; RTO-MP-050), Col. R. Alward in his keynote address said:

"Cyberspace is a new battlespace which subsumes the physical. Unfortunately we are not yet able to adequately *visualis*e it, hence are unable to fight in it.

Consider the impact of the topographical map. It allowed military commanders to expand their battlespace; it allowed them to visualise their space. We have gone from a commander standing on a hilltop to better visualise their space to the use rich 3D topographical maps. Remotely piloted vehicles give commanders a bird's eye view of his space. Indeed it is too much information, too granular. I would suggest the topo maps have allowed military commanders to deliver any amount of kinetic energy (bomb) to any place on earth with great accuracy (through the window).

We need a similar visualisation tool to operate in 'cyberspace'."

"Cyberspace" consists of the operations of computers and of the traffic on the linkage networks among them. It was not the first recommendation heard by IST-021/RTG-007 and its predecessors that the visualisation of networks should be studied. The issue of how to present networks of relationships was posed in connection with peacekeeping operations as early as the IST-020/RWS-002 Workshop in 2000, when it was put as follows (M. M. Taylor, in RTO-MP-050 AC/323(IST-020)TP/10:

*Many of the problems […] have in common that the key requirement is for the user to visualise **relationships**.*

The commander may need to visualise who talks to whom, who can authorise what kinds of action on the part of the belligerents or the NGOs, how civilian refugees relate to each other – who is likely to help whom, and who is likely to harm whom. Are there family relationships to consider?

The field officer may need to visualise the political relationships relating to an unexpected roadblock or an opportunistic encounter with a suspected war criminal. […]

Consider the problem of a field officer unexpectedly seeing a person who might be a war crimes suspect. Before making an arrest, the officer needs to know the ramifications of making or not making the arrest at that moment. The available display is on a hand-held device, but it must display relationships such as the political tensions and support systems around the person, the position of the person within the belligerent organization, and so forth. How? We do not know. Research is needed.

In peacekeeping operations, the commanders and soldiers on the ground need to be aware of the networks of social and political relationships among the individuals and groups with whom they must interact. In anti-terrorist intelligence gathering, the networks of message traffic are a major source of information. The same has long been true in more conventional warfare; one of the covers for the World War II D-Day landings was the construction of a fictitious army in Kent that existed only as a realistic network of radio message traffic among non-existent units that was intended to be intercepted by the Germans. Cyberspace is an increasingly important potential battlespace, but it is not the only one that requires humans to visualise dynamic networks.

To represent the situation of a large number of entities in a geographic battlespace in a way that the commander can visualise the probable course of events is not easy, but it is appreciably easier than to represent the dynamic behaviour of a complex linked network. The cyberspace of which Col. Alward spoke involves the linkages among several million different computers, and what matters is not merely the fact of their linkage, but the nature of the traffic among ill-defined subsets of them.

At the 2000 workshop in Canada, W. Cunningham, a senior advisor to the US military, had discussed the visualisation requirements of intercontinental movements of troops. This problem, too, involves networks, but these networks are of time and resource availability. Troops cannot be loaded onto aircraft if the aircraft are on the other side of the ocean, or if there are no buses or trains to take them to the airfield. When resources are in place is as important as whether they are in principle available.

Infrastructure protection was another topic that requires consideration of different kinds of network was brought up at several of the workshops:. The networks of concern are not only the physical networks such as transportation networks, or electricity and water supply networks, but also the causal and risk-related effects of one network on another. Although the causal factors are real, the networks are conceptual ones that exist only in the minds of the people concerned. For example, in the absence of gravity-flow water distribution, loss of electric power would lead to loss of water supply, which in its turn could lead a city to be highly susceptible to a major fire. At the same time, food in refrigerators would spoil, leading to the possibility of substantial food shortages and possible civil unrest. The temporal dynamics of such effects on the infrastructure demand representation in such a way that the civil protection authorities are able to visualise the probable consequences of any untoward event affecting the civic infrastructure. The problem is a difficult one of network representation.

5.5.2 Depiction of Networks

Several presenters at the Halden workshop offered approaches to representing networks of different kinds. One of the issues, not unique to the representation of networks, is that most networks of interest have far more links than can comfortably be viewed on the screen at the same time. In representations of masses of objects in a field, there are several solutions, the two most basic being to show the whole data set at varying levels of resolution, or to show part of the dataset and allow the user the means to navigate within it to see the whole dataset over the course of the exploration. The so-called “fisheye” representation combines both these approaches into a single interactive display, in which a focal part of the dataset is seen at high resolution, the resolution diminishing with “distance” from the focus so that the whole dataset is accommodated within the single display. The user then has interactive control over the location of the focus, and perhaps of the resolution within the focal area.

When moving among resolutions in the display of a network, problems of aggregation and differentiation arise that are not encountered in displaying objects. Connectivity is a critical aspect of a network. If A is connected to B at the highest resolution, then if the resolution of a display is such that A and B appear, their connectivity must be retained. Likewise, if A and B are not connected, they should not appear to be connected, even at the most coarse resolution. This problem is compounded if the links among nodes have an associated property such as capacity, as might be the case for a road network. If A and B are towns connected only by a goat track over a mountain pass, they are indeed connected, but they should not be shown as connected on a normal road map. On the other hand, for military manoeuvre, A and B may not be connected by road, but the terrain between them may be well suited for the available vehicles, and they should be shown as connected. Aggregation and differentiation issues cannot be solved by fiat, based only on the data to be displayed; they must depend also on the requirements of the user of the display.

Despite this stricture, many basic reductions can be done in some kinds of displays. For example, a road map shown at continental scale will very seldom show a dirt road. It could, if that dirt road was the only pass over a mountain range within several hundred kilometres, but that would be abnormal. As a first pass, road maps at large scales would be likely to show only the most major highways, whereas maps at scales showing city blocks might include even footpaths. At the Toronto Workshop (2004), Bjørke presented a method based on information theory and topology to provide a more effective way of presenting road maps at different levels of detail. In Bjørke's approach, the dirt road that was the only pass over a mountain range would be retained at all scales of display, because it addresses the issue of connectivity retention during network aggregation.

5.5.3 Networks, Data Types, and Displays

At the 2004 Toronto Workshop (IST-043/RWS-006), one of the syndicates defined a network as "*an array of nodes and links that exchange "stuff" on containers under a certain protocol and fallowing a determined path*", and produced a set of network characteristics that a user might want to visualise (Table 5-1):

Table 5-1: Some Attributes of Networks Relevant to Visualisation

• Constraints
  • Rules
• Nodes
  • Location
  • Node Type
  • Symbology
  • Open/Closed
  • Node ID
  • Input/Output Property
• Links
  • Capacity
  • Weight
  • Strength
  • Direction
  • Availability
  • Type of Traffic
  • Location
  • Route
  • Identification (Unique)
  • The Medium
• Thresholds and Changes
	• Traffic
	  • Flow
	  • Path
	  • Routing
	• Topology
	  • Symmetric/Asymmetric
	  • Boundaries
	  • Layout
	  • Hierarchical
	  • Tree
	  • Topology Evolution in Time and Space
	• Logical/Physical
	• Redundancy
	• Partially Connected Graph
	• Protocols
	• Networks Interconnections
	• Location
	• Open/Closed
	• Layers
	• Hierarchy

This list was considered to be a first "straw-man" proposal. To it, one may add variables of network dynamics and some more about the network topology. For example, traffic over a link may take appreciable time, allowing for feedback loops that could create dynamic instability. The topology might be scale-invariant (small-world), clustered, random, or in other ways mathematically defined. Links might be dynamically variable, traffic might be broadcast (simulating a great number of individual links) or with predefined source and receiver (over a specified link). Allowing for broadcast phenomena such as radio transmission or airborne disease vectors would extend the syndicate's definition of a network, by eliminating the final phrase "*and following a determined path*".

A user concerned with a particular application is likely to be concerned with only a few of these variables, but a "Reference Model" has to allow for all of them. The syndicate showed in three example domains how the different variables might differ in importance across application areas. The large table of attributes (Table 5-1) can be substantially reduced for a specific application, but some of them will be important for most uses.

The topology of a network may well affect the way it should be presented for visualisation. If, for example, the distribution of links is random over the nodes, there is no preferred viewpoint, and methods appropriate to geo-referenced data may be appropriate, such as interactive fisheye views. On the other hand, if the network is tree-structured, views that concentrate on a few branching levels within single major branches may be more useful. A scale-free network, on the other hand, may be suited to interactive 3D presentation (as was shown for the depiction of inheritance and message-passing structures in a large software system at one of the early NX workshops).

5.6 INTERACTION

At every NATO and NX workshop, speakers have made the point that for their specific purposes, easy "transparent" interaction greatly improved the user's ability to make sense of a display. At a very early NX meeting, before the initiation of IST-021, one talker demonstrated that using a dynamic 3D display of the inheritance and message-passing network structure of a large software project, interaction allowed the user a large improvement over passive observation of the same kinds of movements of the 3D display. It is not feasible to assess the effectiveness of a display system without considering how the user might interact with it.

Interaction has several different facets, but they come down to two main capabilities: the ability to select from the mass of data those that are to be displayed, and the ability to influence how the selected data are displayed. The user's requirements always determine what should be displayed. The VisTG Reference Model (Figure 5-2) makes this dependence explicit. It also makes explicit the need for interactivity more generally in the construction of systems intended for visualisation.

Selection might take the form of navigation, using a limited field of view, through a dataspace that has some located properties, as would be the case for geospatial data, or, more abstractly, for data located by distance from specified nodes in a network. On the other hand, it might involve algorithmic selection, as, for example, if the user wanted to see railway lines and pipelines, but not roads and electric power lines on a map. Visualisation is hindered if the user must pay attention to the process of selecting the data, thus detracting from the attention that is available for understanding the implications of the data.

Visualisation is also hindered if the user cannot easily relate what is presented in one view or at one moment with what is presented in another view or at another moment. Hollands and Keillor (Chapter 4) examine this in the context of dynamically changing viewpoints in the same scene from exocentric to egocentric and from

2D to 3D, though they consider only changes imposed on the viewer, rather than changes interactively controlled by the viewer. Smestad (Annex A) analyzes the relationships among simultaneous views using concepts from information theory. The smaller the amount of information required to relate one view to another, the more effectively can the user visualise the import of the representations.

The benefits of allowing the user to control the change of view or of viewpoint in a “natural” way are demonstrated theoretically, experimentally, and in the case of networks representing a large software system, practically.

5.7 CONCLUSIONS

A series of workshops have been conducted under the aegis or patronage of IST-021 and its predecessors. They started with annual workshops run by the Network of Experts on alternate sides of the Atlantic. Those workshops proved valuable in clarifying many issues relating to visualisation in areas of interest to the IST group. IST-013/RTG-002 (the predecessor of IST-021/RTG-007) then organized a NATO workshop on visualisation in Quebec, Canada, in 2000, to which operational military personnel were invited. The interactions among scientists, developers, and users proved fruitful, but feedback after the workshop suggested that each community had had some difficulty interpreting the language used by the others, and each felt that the meeting had been overweighted to the interests of the others.

IST-021/RTG-007 organized two workshops, the first (IST-036/RWS-005) in 2002 in Halden, Norway, and the second (IST-043/RWS-006) in Toronto, Canada. These workshops used both plenary presentations and small working groups (syndicates). Members of all three communities (military, developer, and research) were included in each syndicate. These syndicates addressed specific topics in depth, and provided a forum in which the members of the different communities could improve their understanding of each other as well as of the topic. The same approach was used at the N-X meeting in 2003, which also followed up on some of the work done at the 2002 Halden workshop.

All these workshops brought out two major themes:

- Research is needed on the presentation of networks for visualisation; and
- The user’s interaction with the presentation is important in developing effective visualisation.

The VisTG Reference Model served as a framework within which many of the syndicates at each of the later workshops worked, and even though it seems insufficient in itself to provide the means for evaluating visualisation systems, it seems a natural way to think about the problems.

5.8 REFERENCES

[1] RTO Technical Report TR-030, Visualisation of Massive Military Datasets: Human Factors, Applications, and Technologies, Paris, NATO, 2001.

[2] RTO Meeting proceedings RTO-MP-050, Multimedia Visualisation of Massive Military Datasets, Paris, NATO, 2002.

[3] RTO Meeting Proceedings RTO-MP-105, Massive Military Data Fusion and Visualisation: Users talk with Developers, Paris, NATO, 2004.

Chapter 6 – CONCLUSIONS AND RECOMMENDATIONS

The objective of IST-021/RTG-007 was to gain visualisation knowledge in common which will help facilitate the effective development of internationally useful systems. The RTG undertook four major tasks: evaluate and update visualisation systems principles developed by IST-013/RTG-002; deliver a workshop in 2002 on multimedia information visualisation, the results of which helped to guide the subsequent work of the group; propose a lecture series on Multimedia Information Visualisation for delivery in 2005 (this task was cancelled); deliver a workshop in 2004 on Visualisation and the Common Operational Picture.

6.1 EVALUATE VISUALISATION SYSTEMS PRINCIPLES

The visualisation systems principles to be evaluated were of two kinds, a set of guidelines based on information-theoretic concepts, and the VisTG Reference Model which had evolved through the life of IST-013/RTG-002. The former concerned itself primarily with the representations of information on screens, whereas the latter dealt more with the way in which interaction affected visualisation. It was hoped that the two could be effectively combined and used in a coordinated way to provide a framework for the design and evaluation of visualisation systems.

6.1.1 Evaluation Procedure

To test the evaluation procedures, it was considered best that several different systems be used as objects for testing, and that they be tested by people who had not been involved in their development, using the RTG methodologies. Most of the nations participating in the RTG proposed to contribute visualisation systems, but almost all had difficulty in arranging their release to the RTG for testing, despite that the RTG operates under the NATO umbrella with appropriate Memoranda of Understanding. The UK Master Battle Planner, which had been demonstrated to the RTG, was particularly problematic in this respect. It had been promised to the RTG for evaluation, and its release had been approved by DERA. Just before its actual release, DERA was privatized as QinetiQ, and the QinetiQ management rescinded the release permission. In the case of two Slovak systems that were proposed, the difficulty was that the systems and their instructions were only available in the Slovak language, which was not known to the potential testers.

As a consequence of the difficulties in releasing most systems for testing in foreign establishments, it was decided to invert the testing scenarios, and to ask the developers or users in the nations to use the methodologies to assess their own systems, thereby bypassing the issue of releasability while subjecting the methodologies to test. The concept was that the designers or users of the systems might use the methodologies to discover potential improvements in their own systems. In the event, the Norwegian delegate was given a three-day training course on NORCCIS II. He then and on subsequent occasions applied both methodologies to what he had been able to observe. The VisTG Reference Model was applied to the Canadian VITA engine used in the environment of a Web site. Meanwhile, the UK provided an independent evaluation of the "Master Battle Planner" using neither of the RTG methodologies.

6.1.2 Conclusions on Evaluation Methods

The paucity of testing environments makes it difficult to offer a persuasive conclusion as to the usefulness of the RTG methodologies. Certain things did, however, become clear.

- Both the Smestad Guidelines and the VisTG Reference Model need clear documentation as to how to interpret the principles they embody.
- Both require software support, to keep track of the interrelated observations on the different components of the systems under test (depictions in the case of the Smestad guidelines, and levels and parallel loops in the case of the VisTG Reference Model).
- Both systems did identify areas of potential improvement in the systems examined, but those areas had already been noted by users in everyday operation. This agreement may serve to validate the evaluation methods, but do not suggest that they are an advance on the simple experience of users.
- The surface validation provided by the agreement between user experience and the findings from the methods suggests that there might be value in further development of the methods, using supportive software.

6.1.3 Recommendations

The issue of evaluation remains important. Although user experience is often a guide to the deficiencies of an existing system, it can be misleading. "Pretty pictures" may be preferred to duller but more informative ones, and it is almost a human factors article of faith that one does not give the user free control over the colour coding of a display.

Even supposedly expert contractors can produce displays that are handsome but ineffective, as was demonstrated by something that happened to one of the authors of this report (M.M. Taylor). On a visit to a UK Admiralty laboratory, he was shown a submarine tracking display that featured specific detail presented in blue upon a dark background. Although it looked elegant and the users apparently liked it, they had difficulty using this aspect of the display, because blue does not contribute to detail perception. He suggested adding some green to the blue, and was later informed that this one suggestion had produced better results than many thousands of pounds that had been spent with contractors on the display.

The point of this anecdote is to emphasise that evaluations must be objective. The level of user satisfaction must be an element of an objective evaluation, since if users do not like a system, they probably will not use it to its best advantage. However, user preference cannot be the only element. The system exists for a military purpose, and the evaluation must show how well it succeeds in fulfilling that purpose, given the users who have to employ it.

To test a system in realistic situations, while stressing it in all its different modes of operation, is neither feasible nor ethical (since war games seldom offer all the complexity, uncertainty, variety, and psychological-physiological stress of real conflict). Accordingly, evaluation procedures must assess components of the system, and must employ theoretical understanding in order to extrapolate the results to the performance of the system in situations for which it was really intended. Both evaluation methods proposed by IST-021/ RTG-007 were of this kind, and applied primarily to complementary aspects of the visualisation systems under test. If they were to be most useful, their complementary aspects would be integrated into a single evaluation approach.

Individually, both approaches seemed to show promise, despite the small amount of actual testing that was done. Both, however, produce large amounts in interrelated data, which are difficult to use effectively in the form of a paper record or a simple file. To take advantage of their potential, and to provide a proper test, they should be provided with software support. One aspect of this support is in the representation of networks of relationships, which is the topic of IST-059/RTG-025. Another is the possibility of linking into on-line repositories of human factors knowledge that could aid the answering of specific questions asked through the VisTG Reference Model or implied by the application of some of the Smestad guidelines based on information-theoretic principles.

6.2 WORKSHOPS

The 2002 "Halden" Workshop was the first under the aegis of IST-021/RTG-007 to combine plenary "provocation and discussion" sessions with small topic-oriented working groups (syndicates). This format worked well, and was repeated at the following N-X Workshop and at the 2004 "Toronto" Workshop (IST-043/RWS-006). The Plenary sessions cover the major topics of the Workshop and allow all participants to be aware of the ramifications of those issues, whereas the syndicates permit topics of special interest to be investigated in depth, with contributions from all three communities: users, developers, and researchers.

6.2.1 Conclusions

One consistent feature of the workshops was that fewer military users participated than the organizers hoped. Those who did participate made valuable contributions, and their insights provided researchers and developers with an expanded understanding of what was required as opposed to what was currently available to the military. Several of the syndicates either provided new insights into specific issues, or expanded and clarified what was known, in a way understood by the members of all three communities (military, developer, and researcher).

The NX Workshops and the NATO workshops attracted different but overlapping sets of participants, though there was little or no military participation at the NX Workshops. Likewise, the workshops on the opposite sides of the Atlantic attracted overlapping sets of participants. Once the "Plenary and Syndicate" structure had been established, subsequent workshops of both types used it to advantage. Themes that seemed at one workshop to show promise for further development were offered for syndicate consideration at subsequent workshops, provided they were compatible with the overall theme of the late workshop.

Furthermore, the theme of the later workshops were influenced by the findings of an earlier one. The "Common Operating Picture" theme of the Toronto Workshop (2004) was influenced by discussions at the Quebec and Halden workshops that argued for deeper consideration of related problems, as is the network theme both of the forthcoming 2006 "Copenhagen" workshop (IST-063/RWS-010) and of the successor to IST-021 (IST-059/RTG-025). The 2003 NX Workshop theme "Information Visualisation Needs for Intelligence and Counter-Terror" was directly determined by the results of one of the syndicates of the 2002 Halden workshop. The sequence of workshops might therefore be considered as a kind of parallel stream of work, supporting and supported by the overt work of IST-021/RTG-007.

6.2.2 Recommendations

The success of the workshops argues for continuing the principle of alternating NX and NATO workshops, using the established "Plenary and Syndicate" structure, and addressing at each workshop issues highlighted at earlier workshops.

Apart from recommendations about the manner of conducting future workshops, thematic recommendations could be synthesised from the tenor of the work during the workshop. The Halden workshop, for example, provided recommendations that developers should apply human factors knowledge on planned future systems at the time their inputs, outputs and users are identified (this is a perennial recommendation applicable to all military systems), and that NATO should consider the application of visualisation in new areas of military relevance, such as in counter terrorism and coalition operations.

All the workshops reinforced the recommendation that network representation be studied, and that important networks were not only the physical networks such as the Internet or the electrical supply network, but also logical networks such as the different kinds of social relationships that underlie political power in peacekeeping operations. This recommendation is being implemented in the form of IST-059/RTG-025 and the 2006 Copenhagen workshop IST-063/RWS-010). IST-059 Should study generic network visualisation, based on property sets drafted by the Networks Working Group in the Workshop, so as to generate frameworks for networks of potentially very different kinds.

The use of large area high resolution displays should be further explored. Both the table-like display, represented by the Norwegian CODS system, and the wall-like display represented by the US SPAWAR Knowledge Wall have been found useful in practice. Their proponents have offered reasons why they should be valuable, and there are a priori human factors arguments that suggest conditions in which one might be preferred to the other, but little or no research has been done to establish more general conditions under which either form of large-scale display should be used.

At the other end of the size scale, hand-held devices with screens measured in centimetres and storage capabilities measured in gigabytes are becoming readily available. Much research is needed to determine how to select and display situationally relevant information from large dataspaces on such devices, as well as on how to coordinate the local storage with the data available globally to friendly forces.

Issues concerned with very large and with very small displays come together in the following statement quoted from the minutes of IST-021: *Hitherto, the problems of presentation for visualisation have been addressed in a very ad-hoc manner. IST-021 made a start toward systematizing the issues and this effort should be continued, since the fundamentals remain true while the specific applications change.*

Evaluation of presentation technology, including interaction, remains an unsolved problem. Again from the minutes of IST-021: *It is clear that in many cases, user preference is uncorrelated with user performance, but exactly how this happens is not well characterised. It is inappropriate to have standard benchmark and test data sets to compare approaches, since the purposes for which a presentation is being used may vary over a much wider range than the data sets and benchmarks could address.*

Finally, from the NATO viewpoint, the Human Computer Interaction (HCI) implications of differences among nations (or forces) in culture and convention should be carefully studied, with a view to incorporating the effects of such differences into technology intended for NATO deployment early in its development. This point can be addressed from two directions: minimizing the difficulty of interpretation that may be encountered by the nationals of one nation when confronted with technology devised in another, or enhancing the effectiveness of cooperation and collaboration among the forces of different nations (or forces) in multinational (or joint) operations. Since most research in HCI has been done within single institutions, cross-cultural effects have been largely neglected. A NATO RTG designed to address this problem might serve to fill some part of this gap in our knowledge, and in the longer term might assist in multinational joint operations.

Annex A – OUTLINE FOR CONCEPTS AND PRINCIPLES OF EFFECTIVE VISUALISATION

T. Smestad
FFI (Norwegian Defence Research Establishment), Norway

Editor's Note: This Annex consists of a 1993 working paper proposed to IST-021/RTG-007 for testing as a method of evaluating the display properties of a visualisation system. Its use should complement the use of the VisTG Reference Model. Both were tried in the evaluation of the Norwegian NORCCIS II system.

ABSTRACT

Effective visualisation here means to transfer a number of ideas and complex messages in a short time by means of graphics. Thirteen guidelines for such graphics are formulated, partly based on concepts and principles in perception, cognition, and information theory. The guidelines involve visualisation concepts and principles that have evolved over many years starting in the area of technical automation. Five examples are judged in this paper according to the presented framework, three from the first book of Edward Tufte. There seem to be good experiences with the guidelines, but rigorous experiments to measure the effectiveness of graphics are not performed. Such experiments might lead to establishment of a theory of effective visualisation.

A.1 INTRODUCTION

Graphics are currently regarded as a powerful communication channel for the transfer of complex messages and understanding. The use of graphics is increasing, especially due to affordable powerful computers, high-resolution graphical screens and printers, and associated software. However, there is no established knowledge of how to *visualis*e efficiently (design graphics that make viewers understand a lot in a short time). This paper is a summary of possibly relevant concepts and principles that the author has gradually formulated since 1978. This paper is the result of an engagement by the OECD Halden Reactor Project at Institutt for Energiteknikk in Halden to make a brief summary of today's status. Computer graphics in control rooms is one of their concerns; questions from this area are therefore mentioned.

Efficient use of two mutually supporting documentation forms in a software project in 1978 started the evolution of the current ideas. The needs for understanding and documenting technical issues in programming of technical systems have been an important nutrition since then. This problem may be formulated as a static one: given a complex message, how to diagram the message in order to convey it in a short time. Note that the problem formulation of Man-Machine graphics in control rooms is wider: there is no author with a message, and the understanding (of the current situation) may possibly first be obtained over time and by observing responses to induced control actions. Hopefully, the concepts and guidelines presented here are still useful.

The two books by Edward Tufte: "The Display of Quantitative Information" [1] and "Envisioning Information" [2], are probably the best ones written on this subject. The concepts and principles presented in this paper may be seen as other ways of explaining why the good examples in Tufte's books are efficient visualisations. A theory of visualisation, if possible to establish, will probably involve the areas of information theory, cognition, and graphical design. This paper may be seen as an attempt of selecting proper parts from these areas in addition to own concepts for such a purpose. Earlier versions of concepts and principles are presented in the references [3 – 7]. Some of them are used in a Dr. Ing. Thesis [8].

Over the years I have experienced that the presented concepts and principles have inspired the design of more efficient illustrations, both for myself and for other. My hope is that this paper also will do so. In addition, I hope that the paper may bring about a discussion and refinement of the concepts and guidelines, and strengthen the readers' general interest for visualisation. Hopefully this may lead to the design of experiments that may verify guidelines, thus starting the establishment of a theory of effective visualisation.

After this introduction, the paper continues (in Section 2) with reviewing concepts and assumptions from perception, cognition and information theory. Then (in Section 3) new concepts for classifying and describing graphics are formulated. Concepts from all these areas are then summarized in a vocabulary list. Guidelines for effective visualisation are then formulated (in Section 4) on the basis of the presented concepts and assumptions. The deductions are somewhat presented and discussed together with a sequence of questions to answer in order to make efficient visualisation. A few examples are discussed in the presented framework (in Section 5). The establishment of a theory of visualisation is briefly commented (in Section 6) before the summary and conclusion.

A.2 TOPICS FROM PERCEPTION COGNITION AND INFORMATION THEORY

Viewing and understanding messages represented in graphics involve complex processes of perception and cognition. Concepts and assumptions from these fields are a basis for some of the discussion in Section 5, so also with information theory. Key issues from these fields may only be properly treated in textbooks, the presentation here is only a brief review.

A regarded illustration is mapped to the retina where contours, lines, and surfaces probably are identified by a few levels of neurones behind the retina, see [9]. Higher level structures of the illustration (e.g. objects, graphs, tables) are identified a bit further down the processing line where these are compared with similar previously stored structures. The brain probably sorts out, relates, and compares structures and "units" on several levels; see [10]. We are able to simultaneously treat 3-7 such independent "items" on the conscious level. The viewer may normally decide which set of 3-7 items to focus on. These may vary from covering the whole illustration to just cover a little part of it. The viewer's interest in terms of question(s) and uncertainty may eventually be resolved by one or several comparisons of stored images and previous knowledge. The power of this decoding and interpretation process seems to be more dependent on an enormous "database" of stored images and knowledge than on the processing capacity. The information of the message (see definitions in Section 3) is extracted as "differences" between the observed and the expected (or possible) on several levels. In terms of information, the whole process "reduces" the image-information of possibly some tens of millions bits of an illustration to only a few bits of the original question(s) or uncertainty. These concepts and principles are illustrated in Figure A-1.

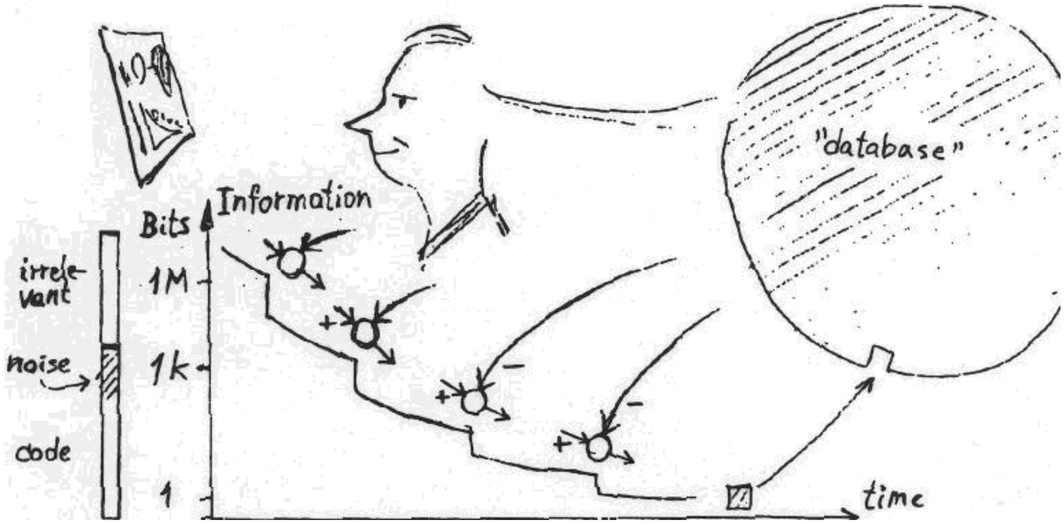

Figure A-1: The Reduction of the Image Information of an Illustration to that of the Message.

It seems that knowledge of complex issues often means that the occurrences of details (of objects, relations, sequences, etc.) are sorted and organized in abstract "scripts". These are much like the class-hierarchy and object-hierarchy used for system analysis according to the object oriented paradigm, see [11] for the latter. Many feel their knowledge in many cases is represented and organized by a sort of inner view. This may possibly occasionally be represented in a fairly one-to-one fashion by graphics. This inner view may be strongly influenced by observed graphics when acquiring the specific knowledge.

Below are statements summarizing the assumptions and principles mentioned. The term "visual processor" and other terms are defined in Section 3.

a) Our *visual processor* is able to quickly digest high information *images* to *messages* (but not all kinds of high information images!).

b) Our *visual processor* tries to organize a complex visual *image* into a hierarchy of *visual units* and make relations and comparison between the units.

c) Our *visual processor* depends on a huge "database" of visual *images* and knowledge; the information is extracted by series of comparisons (differences) between the observed/interpreted and the associated expected or possible observations/interpretations.

d) Knowledge of complex issues generally involves a linking of the concrete and the abstract; this may take the form of class- and object-hierarchies: the concrete belong to instances of objects, the abstract are general properties of the classes the objects belong to.

e) Our knowledge is sometimes represented and organized by an inner view that has features of an *image*.

f) The conscious part of the brain is able to focus and manipulate 3-7 *visual units* simultaneously.

Information theory has a precise definition of "information" relating to coding and transmission of messages; which is probably applicable to graphics. The information of each received "message' (the choice within a set of possibilities) is defined as the logarithm of the inverse of the probability of the message: logarithm (1 / probability (message)). The unit of this measure of information is "bit" if one uses the logarithm with 2 as the base. One bit of information is accordingly the information of disclosing an uncertainty of 50 percent probability. In computers, one bit means a digit being either a "0" or a "1". This corresponds to the information theory's definition when the coding is done in such an (efficient) way that every digit has an equal probability of 0 or 1. Normally the probability is skewed meaning that each digit carries less than one bit information on an average. The mathematical definition may seem a bit odd. However it may be shown that this is the only formula that satisfies a few axioms, these being intuitively associated with "information". It should be noted that this basic definition of information in a way involves prior knowledge. Getting information means being surprised; the greater the surprise, the more information.

An example of a simple failure situation in a plant can demonstrate the calculation of information. Assume that one knows for sure that one valve among 8 is stuck. Assume that there is no indication nor prior knowledge that points to one valve as more likely to have failed than the other. The identification of the failed valve has accordingly a probability of 1/8. The information content in this message is therefore 3 bits ($\log_2(1/(1/8))$).

Information theory says that messages are transmitted from a sender to a receiver through a communication channel by signals, these being coded representations of the messages. The information content (according to the definition) of the signals has to be at least that of the messages, often it is much higher. This is normally due to inefficient coding, but may be done as a counter-measure against failures in signal transmission. Errors may occur in the channels transmission of signals (noise) and in the decoding process. Signal sensitivity may be defined as the amount of change in the signal when the message is changed. High sensitivity means more reliable message transfer (at the expense of transmission capacity). This because the received (erroneous) signals still will resemblance the signals of the coded message rather than a different message.

The traffic light at a pedestrian walking is an example of a high sensitivity signal. The message (go/walk) has one bit information if the two possibilities have equal probability. The signal has a lot more information being light from either upper or lower light bulb, two colours, and two different images (in Norway it is a person standing or walking). There is a high signal sensitivity since a change in the message changes a lot in the signal. Misinterpreting one of the visual effects (emitting light bulb, colour, and image) does not necessarily means misinterpreting the message (go/walk), thus being a reliable information transmission.

Graphics may be regarded as signals since the graphics is a bearer of the messages to transfer, these being coded in the graphics. These signals have generally high information, here called "image-information", see definitions below. An upper bound of this image-information may be calculated by the number of pixels in the image times the different colours and brightness at each pixels. The information is much less since the probability of the appearance of each pixel is strongly dependent of the neighbours. Even so, the image-information is far greater than the message information.

The difference comes from three sources. The greatest contribution normally comes from the low coding efficiency of graphics; low information messages have to be coded with high information visual cues. Note that the message may still be quickly decoded by the brain, low coding efficiency does not imply inefficient visualisation. The second source is redundancy. Redundancy is additional coding to make the messages easier to interpret or less likely to be misinterpreted (see the example of pedestrian walking). A third, and possibly large source, is irrelevant coding. This may either be messages of no interest or coding of the message which do not enhance the decoding (understanding). The noise in a signal is generally the amount of unintended changes in the signals. Noise may still have a meaning in graphics even if the graphics does not unintentionally change. The communication channel may be extended to include the process of perception and cognition. Noise may then be defined as the amount of signals that is interpreted other ways than intended or not interpreted at all. The division of the image-information in these four groups are illustrated in Figure A-1.

Important statements to summarize points in the text:

a) *Images* may be viewed as *signals* which tr*ansfer messages coded* in them.

b) The *information* of *messages* are normally far less than the *information* of an *image*.

c) The *image-information* may be partitioned in *coding information* and *irrelevant image information*; the latter being *visual effects* that have no *coding* of the *message* or *coding* of irrelevant *message(s)*. The former may involve *redundancy*.

d) *Redundancy* will add to the *signal sensitivity* and increase *message transfer reliability*.

e) A low ratio between the *information* of the *message* and the *coding* does not imply inefficient visualisation.

A.3 SPECIFIC VISUALISATION CONCEPTS

The following brief description of concepts starts with Minard's illustration [1] of the French army's campaign to Moscow in 1812 – 1813, see Figure A-2. The visualisation concepts are here used as if they already were known; they are described in the end of this Section. Important concepts are summarized in Figure A-3 which should be consulted while reading. The numbers in parenthesis in the following section refers to the message groups and the visual effects in Figure A-3.

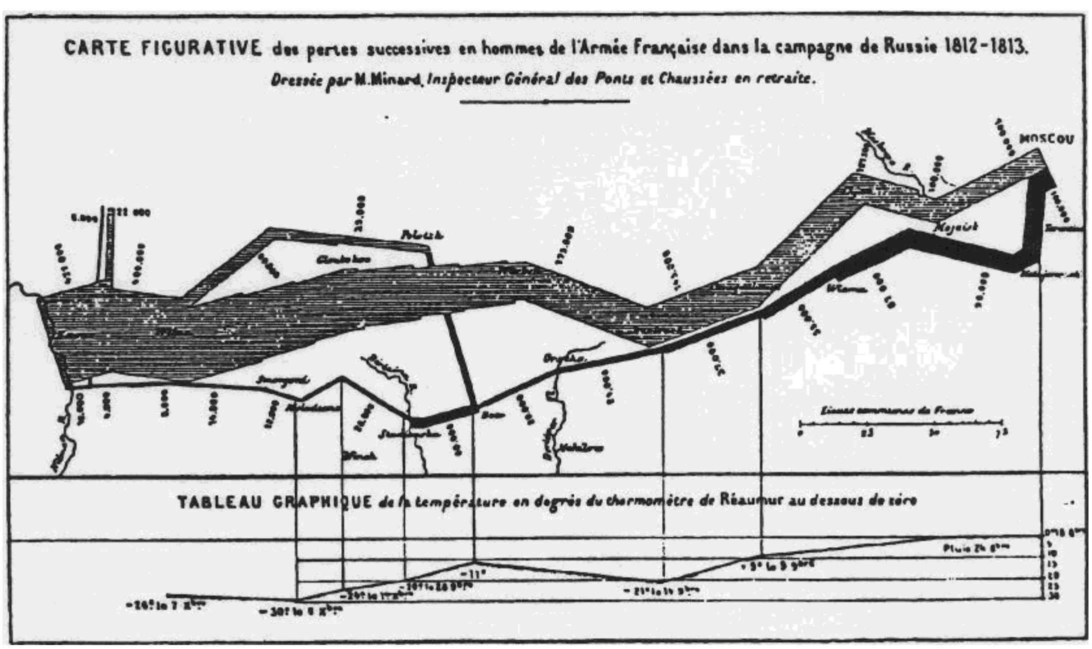

Figure A-2: Minard's Illustration of Napoleon's Campaign to Moscow 1812 – 1813.

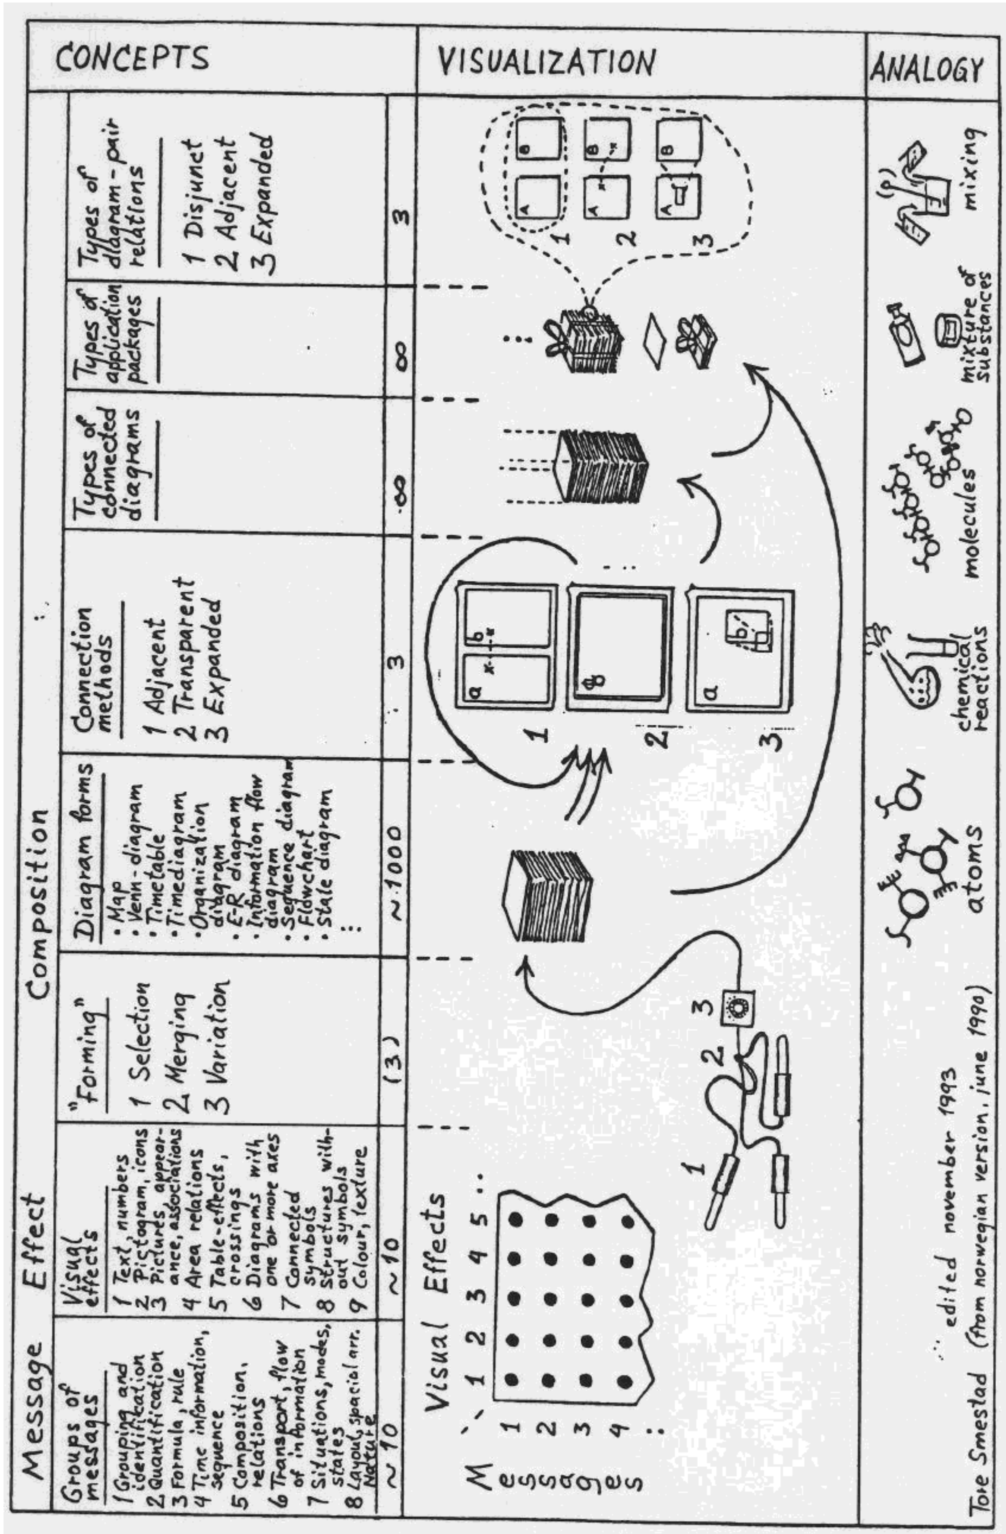

Figure A-3: A Visual Summary of Some Specific Visualisation Concept.

Minard's illustration is a good example for this purpose since it uses 5 types of messages, 5 types of visual effects, 4 diagram forms, and the three connection methods. (Some authors regard this illustration as the best ever made!) The French army's march from the river Niemen to Moscow is visualised with a "band-graph' that codes the quantity as the width (in a logarithmic scale), the direction of the march by texture. Message groups in this diagram form: quantity (2), transport (6), situation, mode (7). Visual effects in it: structure without symbols (8), colour/texture (9). The size of the army at specific points is described with numbers. This may be viewed as an expanding connection using text/numbers as a diagram form. Message group in the diagram form: quantity (2); visual effect in it: text/numbers (1). The "band-graph" is transparently connected with a chart showing the main rivers and name of places. Message group in this diagram form: layout/spatial arrangement (8); visual effect in it: picture/appearance (3). The shrinking of the army on the retreat was highly due to the low temperature, which is diagrammed below. Message groups in this diagram form: quantity (2), time information (4); visual effect in it: diagram (6). The temperature diagram is adjacently connected to the (returning) "band-graph" in that the time-axis and the charts east-west axis coincide. Note that this connection enables the identification of the timing of the retreat since the temperature curve is labelled with the time. Find a summary of the used diagram forms, their size, messages, visual effects, and connections in Table A-1.

Table A-1: Diagram Forms and Connection Structure of the Examples

Diagram form	Area	#	Mg.	V.e.
Figure A-1:				
Person viewing an illustration	30%	1	6,8	3
Illustration of size of the brains “database”	20%	1	2	4
Diagram of reduction of the information	40%	1	2,4,6	6
Illustration of the many comparisons	3%	4	3,6	7
Bar with groups of information	2%	1	1	4
Figure A-2:				
Chart of part of Europe	40%	1	8	3
“Band graph” with texture	50%	1	2,6,7	8,9
Numbers	1/5%	25	2	1
Temperature diagram	30%	1	2,4	6
Figure A-3:				
Table	100%	1	1	5
Texts	30%	8	1	1
“Flow” from messages to illustrations	25%	1	6	7
Illustrations of elements in the “flow”	15%	10	1	3
Illustrations of analogies	5%	5	1	3
Figure A-4:				
“Diagram” of air pollutants and time	85%	1	1,4	5
Diagram of concentrations at places	7%	12	2	6
Chart over the Los Angeles area	15%	1	8	3
Figure A-5:				
Diagram of height/depth over the year	100%	1	2.4	6
Drawing of the insect in nature	100%	1	2,8	3
Drawing of blades at breeding time	15%	1	8	3

ADJ. TRA. EXP.

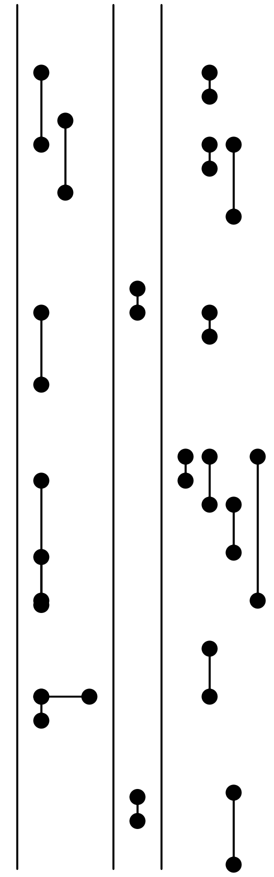

Area: % of the image area “covered”

#: Instances of the diagram form

M.g. : Message group

V.e. : Visual effect

ADJ.: Adjacent connection

TRA. : Transparent connection

EXP. : Expanding connection

Above, the terms diagram form, message group, visual effect, adjacent-, transparent-, and expanding connection were used. These terms are loosely defined in the end of this section, and they are illustrated in Figure A-3. The specific grouping of “message groups” and “visual effects” in the figure are not essential, other possibilities exist. (These ones may be biased from the area of software and industrial automation where they emerged.) The important point is that an illustration carries messages (that may be grouped) which are to be transferred to the observer by representing them (coding them) by visual effects (that may also be grouped). This division of messages (the information to transfer) and visual representation (the signals) is assumed important.

The three connection methods may be regarded as techniques in placing more information in the same image. Even though these may be seen as layout-techniques, the success of a connection (the ease of understanding the messages from each "component" and their relation) is highly dependent on the viewer's interpretation. This means the connections both have a graphical and a semantic aspect. The connection methods should not be regarded as precisely defined, intermediate connections may be found.

In Figure A-3 the ability of a diagram to form connections is depicted as "hooks" on the atoms, which may be viewed as their "valence". The diagrams must have some elements in common in order to be connected. These may be seen as the elements in an n-tupple. The situation of the French army (Figure A-2) at each point in time may, among a lot other variables, be characterised by: geographic position, direction of march, weather condition (temperature, wind, rain/snow), total number of individuals and their grouping (regular soldiers, artillery, cavalry, etc.), equipment, and supply of food. A point (cross-section) of the "band-graph" in a north-west map will represent 3 elements in the n-tupple: position, direction, and (roughly) the number. Other elements of the army's total n-tupple at that moment could be represented by other diagram forms. Only the number is described; this is more accurately represented by digits than by the graph. The (whole of the) "diagram form" of digits represents a single point of the graph, therefore classified as an expanding connection.

The rest of this section is a kind of glossary of important concepts and terms used in the paper:

coding:

representing the *message* by visual means using a *coding scheme* that the viewer understands

c. scheme:
a specific way of representing a *message* visually to convey it, c.s. produces the *visual effect*

connected:

1) property of an illustration consisting of at least two instances of *diagram forms* being a part of a *connection*
2) property of an instance of a *diagram form* in an illustration being the part of a *connection*

connection:

linking between two instances of *diagram forms* by utilizing the representation of one or several common elements or n-tupple(s); there are three types of *connections: adjacent c., transparent c.*, and *expanding c.*)

c. method:
one of the three following way of linking instances of *diagram forms*

adjacent c.:
where the two instances of *diagram forms* occupy different areas of the illustration and one or more elements on one corresponds to one or more elements in the other

transparent c.:
where the two instances of *diagram forms* occupy the same area of the illustration and what the positions represent in one instance correspond to what the positions represent in the other

expanding c.:
where (the whole of) one instance is related to a single part in the other; the two instances may or may not occupy the same space of the illustration

(Comments to the connection methods described:

an illustration with a single instance of a diagram form is of course not connected; having two or more instances, the illustration may have one or several connections. The connections may form a complex structure since a diagram form can be connected to several other and a single connection or a structure of a connection may be further connected as a unit.)

connection key:

independent unit(s) or instance(s) of element(s) in n-tupple(s) being represented in a *diagram form* enabling the *diagram form* to be *connected* to other instances of *diagram forms*

diagram:

d. form:

a class of diagrams characterized by the types of *message(s)* and *visual effect(s)* used

(Comment to this concept:

according to this definition a table is not a diagram form without stating the *group(s) of message(s)* it represents)

d. instance:

one specific item (instance) of a diagram form or a structure of connected diagram forms

effective visualisation:

illustrations and graphics resulting in a high *message transfer rate*

image:

1) picture in terms of its least graphical elements (i.e. pixels, lines, colours)
2) the spatial area containing the graphical elements (the complete illustration)

mental i.:

an inner view representing knowledge having characteristics of an *image*

information:

coding i. (in an illustration)**:**

the *image information* of all the *visual effects* representing *messages* or *coding schemes*

image i. (in an illustration)**:**

the *theoretical information* of the *image*

irrelevant (image) i.:

the *image information* of *visual effects* irrelevant for interpreting the *messages* or *visual effects* of *irrelevant messages*

message i. (in an illustration)**:**

the *theoretical i.* of the *message*

redundant i. (in an illustration)**:**

the *image information* of *coding* not strictly necessary to convey *the messages*, see *redundancy*

theoretical i. (in an illustration)**:**

the quantity calculated according to the mathematical definition in the theory:

= log2 (1 / probability (subject))

message(s):
the underlying idea(s) or story to convey by the illustration

m. group:
messages sorted according to some classification (here a classification of 9 groups is made)

m. transfer:
the completion of the reception of the ideas being represented in the graphics

m. transfer rate:
the amount of messages transferred pr unit of time

m transfer reliability:
the probability that the intended message is received correctly

noise (of an illustration)**:**
1) (an instance) *visual coding* that are misinterpreted or has delayed the interpretation
2) (characterization) *visual coding* that are likely to be misinterpreted or delay the interpretation

redundancy (in an image)**:**
part of the *image* that are not strictly necessary for interpreting the message, consists of additional *coding* by use of more than one *coding scheme* or clarifying of *coding schemes*

signal:
carrier of the *messages* through a communication channel, here: an *image* that contains *message(s)*

s. sensitivity (of a diagram form)**:**
the amount of perceived change in the image as a result of a change in the message

valence:
(here) a property of a *diagram form* or illustration: the ability to make *connections* to other *diagram forms* by the representation of identifiable elements or n-tupple(s); corresponds to the number of *connection keys*

visual:

v. processor:
a term used for the joint calculating features of the eye and brain making it possible to digest the *image* projected on the retina to meaningful *messages* (the visual processor cannot do its job without an enormous database of stored *images* and knowledge)

v. cue:
features of the *image* that may carry or help decoding a *message*, may be equivalent to *v. effect*

v. effect:
feature of the *image* that may carry some *message*; here these are sorted into 8 groups

v. structure:
the structure of relations between the *visual units* and associated *connection keys* in an *image* as intended by the designer

v. unit:
the whole or a portion of an *image* that is treated as a unit in some respect; examples are *images* representing a physical object, a graph, a variable or a group of such "units"

A.4 DEDUCIBLE GUIDELINES FOR EFFECTIVE VISUALISATION?

This section states guidelines and presents deduction relations from assumptions and principles stated earlier. Guidelines are tried formulated for experimental support. A guideline is here assumed supported if experiments show a positive correlation between its degree of being met and the message transfer capacity. The terms “degree of being met” and “transfer capacity” have to be quantified in order to do such experiments. This topic is beyond the scope of this paper. There are several important questions not explicitly treated in the presented guidelines. Some are mentioned in the discussion at the end of this section as warnings and as inspirations to further thought.

Below is presented 13 guidelines for effective visualisation. There will of course be a lot of other guidelines aiming at other levels of detail and using other terms. One example is: “an image should be comfortable to look at having proper contrast, adequate resolution, soft colours, and having well balanced spaces”. Such guidelines are seen as supplementary to the ones here.

1) Place a let of *messages* into single *images*.
2) Make illustrations with high *image-information* and a low ratio between the *irrelevant image information* and that of the *coding-information*.
3) Partition an *image* into gradually smaller parts.
4) Arrange the *visual structure* so that 3-7 visual *units* may be naturally focused at a time.
5) Select *diagram forms* with a high *signal sensitivity* of the *message(s)* to transmit over the actual range of *message(s)*.
6) Use obvious *coding schemes*; don’t change them, clarify them explicitly if not obvious.
7) Select *diagram forms* and *coding* that are likely to resemblance human’s *mental image*.
8) Make illustrations with a *visual structure* and content according to prior knowledge.
9) Use several *diagram forms* in the same image.
10) Select *diagram forms* representing different levels of abstraction in the same *image*.
11) Pair many instances of n-tupple elements of the selected *diagram instances* by the *connection methods*.
12) Use transparent and *adjacent connection* to minimize the position-uncertainty of *visual units*.
13) If the illustration does not contain all *messages*, select *diagram forms* with *connection keys* to (potentially) supplementary illustrations (that is: furnish *valence* to the illustrations).

Guideline 1) is motivated by the very definition of efficient visualisation and assumption a). The last part of guideline 2) is obvious since irrelevant coding or messages will slow the interpretation process. High image information then means high coding information, giving room for effective coding schemes. Note that there is no claim that a low ratio between the image’s coding- and the image’s message-information will increase the efficiency (the ratio will always be much greater than one). The image-information of an icon is normally much higher than that of the corresponding function name, but we understand it quicker.

Guideline 3) is a simple way of obtaining 4). Guideline 4) is motivated by assumption b) and e). Guideline 5) is another way of stating one should select a diagram form that clearly shows the message and discriminates between variations of it. Guideline 6) is obvious; if the coding scheme is not obvious, it will take longer time to understand the image if possible at all. A change normally confuses and may lead to misinterpretation.

Guideline 7) and 8) is motivated by assumption c) and that information means surprise. It seems reasonably that the message transfer rate is increased if the comparisons are based on a well-known base, see Figure A-2 Guideline 9) is especially aimed at complex messages with several aspects. Then a single diagram form is unlikely to satisfy guideline 5), therefore several diagram forms. Several diagram forms may also contribute to the guidelines 1), 2), and 4).

Guideline 10) is motivated by assumption d). Guideline 11) asks for a sort of visual network of the n-tupples of the message. Making the relations between the n-tupples more apparent will add to the information and coding of the image, thus contributing to guidelines 1) and 2). Guideline 12) is obvious since position uncertainty means time to find the elements and make the comparison between them more difficult and time consuming. Guideline 12) means to enable the use of guideline 11) on separate images.

Important questions and possibly objections could be worthwhile to discuss here. Only a few topics is treated below.

Guidelines 1) and 2) state the images should be of high information density. Generally, there will be a tradeoff between the information content of each image and the number of images. The information density where one is better off splitting the messages on several images, are probably much higher than one believes. The observation time is not explicitly treated. The message transfer rate of a given image will generally be highly dependent on the observation time. The guidelines are formulated under the assumption of a long observation time. This distinction is believed to be important, but is not treated here.

Also, the guidelines are formulated for a static situation: given a message, how to effectively transfer it. This is valid in software documentation, textbooks and user manuals. It is not the case in a control room situation. In a dynamic situation, there is not a static message to transfer. It is more like walking in the dark with a searchlight: which spots to enlighten to verify a safe march? Such a feedback situation will probably influence the set of proper guidelines.

Note that the guidelines generally require knowledge of the observer and the situation (especially no 6, 7, and 8). This means the term “effective visualisation” is no characterization of the illustrations in their own right. If “effective visualisation” is only concerned with a high message transfer rate, what if the messages are ill structured and else unsuited for enlightening the problem domain? (This is claimed to be the case in software engineering before the object oriented methodology was available.) Tufte [1] addresses this question; he states that “graphical excellence” is ...” the well-designed presentation of interesting data – a matter of substance, of statistics, and of design”. The guidelines do not address this point, with a possible exception of no 7), 8), and 10). Prior knowledge and inner representations are believed to be object-oriented and having different levels of abstraction, which is here assumed to be beneficial.

Discrimination and treatment of the concepts of “coding”, “noise” and “irrelevant messages” are not properly done. What Tufte [1] terms “distortion” is a special case of “noise”. To avoid distortion he claims that “The number of information-carrying (variable) dimensions depicted should not exceed the number of dimensions in the data”. Irrelevant messages may also be “noise” as defined here, since they may slow the interpretation. However, the distinction is subtle. In a large map, there is a lot of irrelevant messages for a specific journey, but this does not generally slow down the interpretation of the interesting messages.

Maybe the concept of ergonomics also should play apart. There should be little extra mental effort in understanding graphics, meaning if one gets mentally tired, this ought to come from the mental effort of acquiring the new knowledge at the message-level, not the visual decoding of the graphics. This question may be somewhat addressed in the guidelines 7) and 8), but there may be more to this question.

The production of efficient visualisation will involve a price aspect. Efficient computers, displays, and graphical software are still expensive. It is time-consuming to design and draw illustrations with a high information density, thus costly. The benefit of efficient visualisation will have to be weighted against this cost in a practical situation.

To summarize this section, let us review what is necessary to achieve effective visualisation according to the stated guidelines (note that all of the points have not been explicitly addressed in the paper)

1) Who is to interpret the graphics, what are the questions they have, what do they have of prior knowledge, and what are their internal representations?
2) What are the messages desirable to transmit, which sets of n-tupples are involved?
3) Which diagram forms have the highest signal sensitivity over the range of messages, and do they represent the actual n-tupples suitable for connection?
4) Which set of diagram forms may be connected to constitute an effective single image?
5) How to partition the messages into different images to an efficient "package" of images?

A.5 ANALYSIS AND DISCUSSION OF A FEW EXAMPLES

This section briefly applies the concepts and principles on the diagrams in this paper as examples, two additional illustrations from Tufte's first book [1] are included here. The visual structure of the examples are summarized in Table A-1. Table A-2 summarizes how the examples are judged according to the guidelines.

Table A-2: Judgment of the Guidelines' "Degree of Being Met" in the Presented Examples

GUIDELINES	THE EXAMPLES (FIGURES)				
	A1.1	**A1.2**	**A1.3**	**A1.4**	**A1.5**
1) Place a lot of messages into single images.	1	2	2	3	2
2) Make illustrations with high image-information and low …	1	2	2	3	3
3) Partition an image into gradually smaller parts.	1	2	3	2	0
4) Arrange the visual structure so that 3-7 visual units may …	1	3	3	3	1
5) Select diagram forms with a high signal sensitivity of the …	1	3	2	2	3
6) Use obvious coding schemes; don't change them, clarify ...	2	3	1	3	3
7) Select diagram f. and coding that are likely to resemblance ...	2	3	1	2	3
8) Make illustrations with a ... according to prior knowledge.	1	2	1	2	2
9) Use several diagram forms in the same image.	1	3	3	1	2
10) Select … representing different levels of abstraction in …	3	2	2	1	2
11) Pair many … n-tupple elements … by the conn. methods.	1	2	3	3	3
12) Use transparent and adj. conn. to min. pos.-uncertainty …	–	–	–	3	–
13) If … not contain … diagram forms with connection keys ...	–	–	–	–	–

Scores of "degree of being met":

– : irrelevant guideline
0 : negligible
1 : to some extent
2 : moderate
3 : very well

Figure A-1 is viewed as an adjacent connection of the drawing of a viewer of a diagram and the diagram of the information reduction. The former has m.g. (message groups) : transport (6) and "nature" (8), and v.e. (visual effects) : appearance (3). The latter having m.g. (6) and time (4), and v.e. diagram (6). Distance and time are the connection keys. The diagram of the viewer is expanded by a diagram illustration the vast database of the brain, m.g.: quantification (a coarse one!) (2). v.e.: area relation (4). The diagram of the information reduction is expanded by small diagrams at some points where the retina/brain (the visual computer) is assumed to perform a "subtraction", m.g. formula (3), flow of information (6), v.e.: connected symbols (7). These are also adjacent to the "database". The leftmost part of the diagram is expanded by a column indicating the grouping of information in the regarded image, m.g.: grouping and identification (1), v.e.: area relation (4).

Figure A-2 is already described in Section 3. Figure A-3 might be seen as a table, m.g.: grouping and identification (1), v.e.: table (5). The columns are eight concepts (group of messages, visual effects, ..), the rows are labelled "concepts", "visualisation", and "analogy". The columns of the upper row are expanded by textual descriptions, m.g.: grouping and identification (1), v.e.: text (1). The second row is expanded by a flow diagram, m.g.: flow (6), v.e.: connected symbols (7). The "items" of the flow diagram (the concepts) are expanded by diagrams that associates with them, m.g.: identification (1), v.c.: appearance (3). Note that the columns are not explicitly drawn in this second row, neither in the third showing the analogies from chemistry.

Figure A-4 consists of 12 items of 3-axes diagrams, m.g. quantification (2), v.e.: diagrams (6). However, one may say that the illustration is mainly one big table or rather a diagram with axes time of the day and air pollutants, m.g.: grouping and identification (1) and time (4), v.e,: crossings (5). This big "diagram" serves to uniquely position the 12 small diagrams, which are expansions of the former. The connection keys are here time of day and air pollutant. Adjacently connected to the diagrams is a map over the area, which serves to ease the interpretation of the diagram.

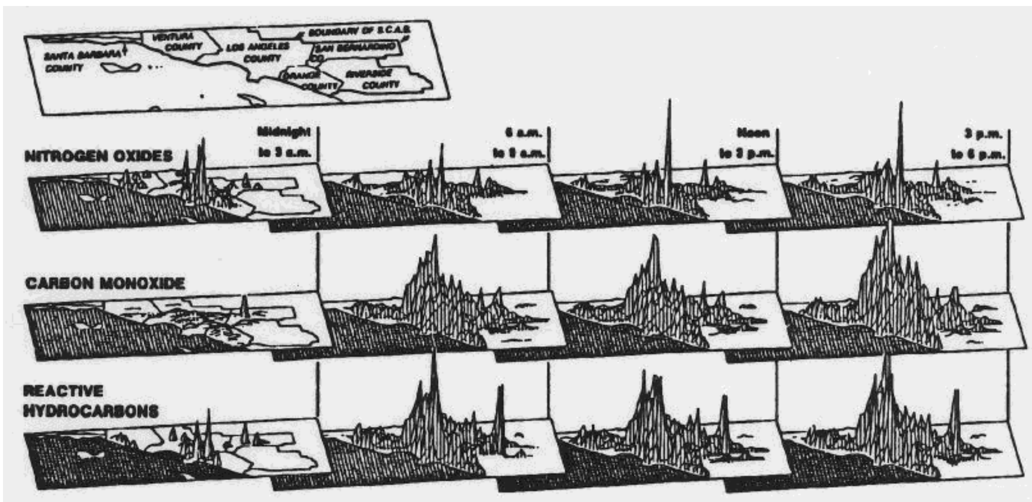

Figure A-4: Illustration of Measured Pollution in the Los Angeles Area.

Figure A-5 has a transparent connection of a sort of natural picture and a partly invisible diagram. The latter shows the height/depth of the insects' location as a function of the time of year, m.g. quantification (2) and time (4), v.e.: diagram (6). The former happens to show the same message (2) due to the same spatial arrangements, but it also shows the appearance of the insect and its surroundings, m.g. layout, nature (8), v.e. picture (3). A picture of the plant at the time of the breeding is expanded, m.g. nature (8), v.e. picture (3).

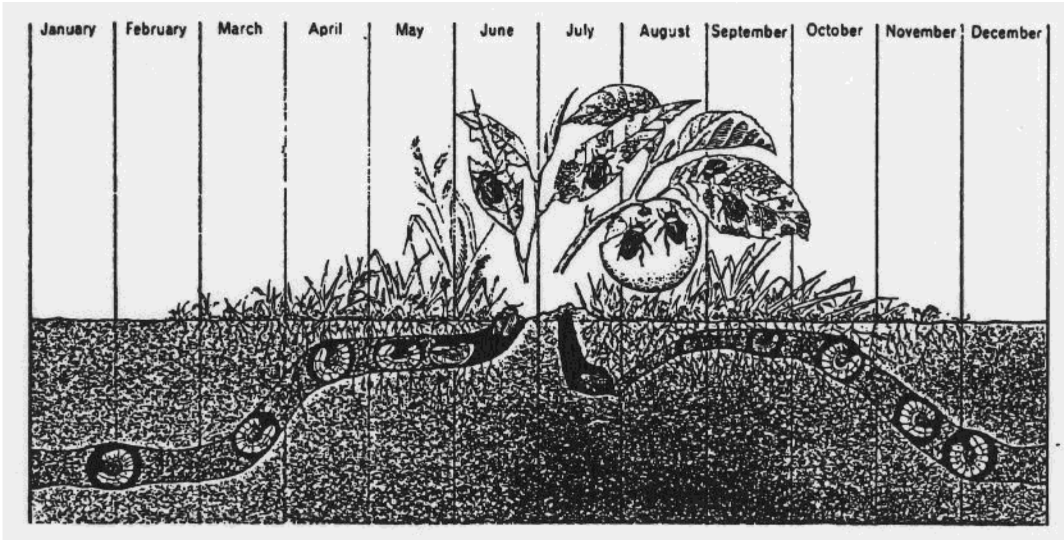

Figure A-5: Illustration of the Life Cycle of the Japanese Beetle.

A.6 STEPS IN ESTABLISHING A THEORY OF EFFECTIVE VISUALISATION

An established theory generally has to contain the following elements:

1) A set of fairly well-defined and self-consistent concepts that describes phenomena of an area;
2) A set of statements (theorems) that may be subject to experiments; and
3) A set of experiments that have the power of falsifying statements in 2, but have not yet done so.

There has to be a fairly common interpretation and agreement of the concepts in 1). An overlap with established theories is regarded as a strength. This means that some of the concepts and statements from established areas may be used and have relevance in the new area, and vice versa. The interpretations of the experiments in 3) have to survive open criticism.

This paper has touched the points 1) and 2); the guidelines are regarded as statements under 2). The acceptance of the visualisation concepts is not established, but it seems that people tend to classify illustrations the same way using the concepts (classifications like those shown in Table A-1). In addition, there seem to be relevant relations to perception, cognition, and information theory. Own experience and other's seem to agree with the guidelines, and the good and bad examples in Tufte's books [1] and [2] also seem to agree with them.

However, no rigorous experiments are done. A next step of establishing a theory of effective visualisation, if possible to establish, is to perform experiments that can test some of the guidelines.

A.7 SUMMARY AND CONCLUSION

This paper has presented thirteen guidelines for effective visualisation (Section 4). These are partly based on concepts and principles from perception, cognition, and information theory, which are stated in summary form

(Section 2). The guidelines are to a large extent formulated with new concepts of visualisation, these and other concepts are summarized in a vocabulary list (Section 3). The relations between the basic assumptions and the guidelines are sketched together with a brief discussion of some important questions.

This edition of the paper is the result of an engagement by the OECD Halden Reactor Project at Institutt for Energiteknikk in Halden. The hope is that these guidelines may inspire the design of new and more suitable control room graphics for Man-Machine Interfaces. The basic questions of this area are wider than those that motivated the development of these ideas, as mentioned in Section 1 and 3.

Even if there are indications that these concepts may be generally accepted and that the guidelines result in effective visualisation, the way to an established theory of effective visualisation may be considerable, if reachable. Experiments that can test some of the guidelines seem to be an important step toward such a theory.

A.8 REFERENCES

[1] Tufte, E. (1983): "*The Visual Display of Quantitative Information*", Graphics Press, Cheshire, Connecticut.

[2] Tufte, E. (1991): "*Envisioning Information*", Graphics Press, Cheshire, Connecticut.

[3] Smestad, T. (1983): "The missing link: Funksjonsspesifikasjon og -beskrivelse av automatiserings-systemer", *Automatisering* nr 8 1983.

[4] Smestad, T. and Andersen, O.K. (1988): "Integrated projection illustrating", *ACM Software Engineering Notes*, January 1988.

[5] Smestad, T. (1989) : "Store muligheter med manipulering av figurer", *Teknisk Ukeblad*, nr 351989.

[6] Smestad, T. (1989): "Programmering av industrielle styresystemer, Mot effektive beskrivelser", *Automatisering*, nr 9 1989.

[7] Smestad, T. (1990): "*Can a new set of visualization related concepts be used to guide the design of man-machine graphics?*", Institutt for Energiteknikk, Halden, February 1990 (not published).

[8] Øgård, O. (1990): "*A situation adaptive presentation system for process control*". Dr. Ing. Thesis, Norwegian Institute of Technology, Trondheim, October 1990.

[9] Davidoff, J. and Concar, D. (1993): "Brain cells made for seeing", *New Scientist*, 10. April 1993.

[10] Albus, J.S. (1991): "Outline for a theory of Intelligence", *IEEE Transactions on Systems, Man, and Cybernetics*, No 3 May/June 1991.

[11] Jacobson, I. (1992): "*Object-Oriented Software Engineering, A Use Case Driven Approach*", ACM Press, Addison-Wesley, 1992.

Appendix 1 – Usable "Visualisation Evaluation" Without System Users!

Tore Smestad, August 23, 2001

A1.1 INTRODUCTION

This is written in preparation for the IST-021/RTG-007 meeting in Denmark October 2001, as an input to the discussions concerning the "visualisation evaluation" we are planning to do; see the initial considerations in point 6.2 in the minutes of meeting from Bonn. Now I generally share Martin Taylor's worries concerning the "evaluation non-progress" as stated in his e-mail of 29.07.01 to us all evaluation volunteers. However, I think the establishment of a workable evaluation method is much more critical at this point than making available the systems to be evaluated.

I think we all share the expectation that we in IST-021/RTG-007 will be able to make meaningful evaluation reports for the developers of the systems solely based on "Our Theory & Method". I think "ordinary usability consultants" have to observe the behaviour and statements of real users in front of the systems – therefore the title of this note. "Our Theory & Method" just has to supply us with this ability, otherwise it must be deemed as non-working. "Our Theory" is currently best formulated in the "HAT-report"; supplements are Justin's Chapters 4 and 5: "Spatial Displays" and "Navigation and Interaction in Real and Virtual Environments", and Tore's 13 guidelines in "Outline for concepts and principles of effective visualisation". "Our Method" is currently formulated in Martin's "Guidelines for evaluating visualisation system designs and specifications". Important elements of a "theory & method" are definitely in place, but I do not find this foundation workable with respect to practical evaluations, mostly because I lack a concrete procedure. (The mentioned references are found on www.vistg.net except Justin's two chapters.)

This note is an attempt to bridge the gap I see between our current foundation and our (lacking?) ability to do "visual evaluations". I see this note as giving practical suggestions towards a "method", not at all as an alternative or in opposition to the mentioned references. An immediate consequence of my suggestions and arguments, is that there is a lot to be done by the "system proposers" before they necessarily supply (key elements of) the system to be evaluated.

It further worries me that I do not see the line between a "visualisation evaluation" and (as I see it) a more general "usability evaluation". I think we should restrict ourselves relative to the latter in order not to "drown" in the vast number of somewhat related problems, and not being mixed with "usability consultants"/ "HCI experts". In a clarifying process we might find our specific "pre"/ "force"/ "speciality" that hopefully can ease and direct our coming efforts of evaluations and of writing a textbook!

A1.2 BASIC SUGGESTIONS FOR WHAT TO DO IN OUR EVALUATIONS

I suggest that we first try to "simulate" the real users. Representative tasks of the users should then be available as "user sequences" (see next section). If the system is not readily available, we should apply "paper prototyping" (see next section) by representative samples of transparent copies of the necessary system displays. This "simulation" may help us to answer many of the questions that Martin poses with reference to details and circumstances the developers understand. To aid us in answering the right questions and when, checklists for doing so should be used.

In doing the "simulation" above, various guidelines might come to mind. These can come from available general HCI guidelines or from other material. We might also be able to create our own guidelines as extracts or syntheses from available material and knowledge.

A last more "off-line" inspection can then be to regard the displays and their interrelations. Here Tore's 13 guidelines can be applied. Also, the list of the various "visual enhancements" described in the database of TTCP AGVis should be applied; which are used in the system and which are not?

A1.3 EVALUATION "TOOLS"

Here is a brief description and comments of the concepts in the previous section. I view them as necessary tools for our evaluations.

A1.3.1 "User Sequences"

In current software system development "use case" seems a central concept, I first read about it in the textbook "Object-Oriented Software Engineering – A Use Case Driven Approach" (by Ivar Jacobson et al, 1992 ACM Press). They say (Chapter 6.4.1): "An instance of an actor does a number of different operations to the system. When a user uses the system, she or he will perform a behaviorally related sequence of transactions in a dialogue with the system. We call such a special sequence a use case." To avoid possible misinterpretations I suggest we use "User sequences" in our case. I think user sequences can be described on somewhat different level of abstraction. The actions at one level might be a complete user sequence at a lower. Also, a more urgent user sequence can interrupt another and being completed before the first is resumed.

A1.3.2 The System Displays

(The system displays should perhaps be viewed as a critical aspect of the system to investigate rather than a "tool", but it seems natural to view "displays" as a sort of tool.) A "display" in this context is very close to a "window" in today's vocabulary. If windows are tiled in a fixed manner, the whole tiled unit is then the "display", possibly the whole screen. A "display" can be physically represented by a transparent for used in "paper prototyping". A "display" is (the sum of) "a diagram form" and "connected diagram" in the 13 guidelines in "Outline for concepts ...".

A1.3.3 "Paper Prototyping"

To enable a rapid prototyping in software development, "paper prototyping" is a simple and low-tech technique; see for example www.infodesign.com.au/prototyping/default.html. I was presented the method as use of overheads with suggested screen-layouts and interactions "simulating" the behavior of a system in response to users. Transparencies work because the "windows" can be freely placed relative to each other, also partly or wholly covering each other. I have barely tried the method, but was impressed by how quickly that shortfalls and problems with a computer system and its HCI not yet built, did appear.

A1.3.4 ("Visualists") Checklists

Pilots use checklists no matter how experienced. Checklists could be an equally important tool for us "visualists" (could it mean "visualisation experts"?). These should guide our detailed questions and judgments of the system in a systematic way. I suppose Martin, Justin and Tore are able to make checkpoints out of the material mentioned in the second section of "Introduction".

I heard from a consultant some years ago that workable checklists never exceed one page. If more point are wanted on it, some other have to be taken out. This might be a guide for the level of detail in a checklist. If a referenced collection of guidelines containing hundreds of guidelines is to be applied (see Section A1.3.6), this can be stated as only one checklist point.

A1.3.5 "Visual Enhancements"

The TTCP AGVis has made a reference model framework called RM-Vis. It shows three main axes, one called the "how-axis" (or "Visualisation Approach"). A description by AGVis states: "Approaches are characterized in terms of the visual representations used (e.g. graphs, charts, maps), visual enhancements (e.g. use of overlays, distortion, animation), interaction (direct manipulation, drag and drop, haptic techniques, etc.), and deployment which includes the computing environment (display device, COTS software) and advanced deployment techniques such as intelligent user support and enterprise integration." I think to remember that "visual enhancements" is a field in their databases for describing a number of systems with advanced visualisation.

A1.3.6 HCI Guidelines

Human Computer Interaction guidelines are natural to use in our evaluation, and we should get hold of some sets. The ones I have happened to see did by no means convince me. We once numbered 1419 in relation to a project. I found them questionable for several reasons, among them a low technical level of detail and a general and non-testable form. An example of the first: "The upper-left-pointing arrow is used for object selection in most windows" (no 13). An example of the latter: "Window layout is logical to operators and appropriate to actions executed" (no 1030).

A1.4 EVALUATION REPORTS

As stated in point 6.2 in the minutes of meeting from Bonn, IST-021/RTG-007 could produce a mini-report for each system. I suggest we start thinking of these in light of the readers (mainly system developers?), their expectations and their benefits of reading. I suggest we define the reports as open to everyone, and get approval from those willing to expose their system to our criticism. They will thus get the reports from the other systems as well. This might be seen in the spirit of free exchange of results within the NATO community, included civilians.

I suggest we try to strike a balance between presenting negative findings/judgments and praise of good solutions. The former should be accompanied with our suggestions of better alternatives and remedies, possibly by giving references to known solutions and existing systems. Comments should be referenced to the checklist points bringing them up and specific guidelines if applicable. In connection with each comment/ item, I suggest we do a sort of evaluation or classification of it, like judgment of the importance of the item for a user, and the "obviousness" of the item. Will a first simple usability test immediately reveal the item, or is "Our Theory & Method" important for finding it? (hoping for a number of such items!!) I think we have to distinguish between experienced users and novices, even between users having picked up different pieces of information in their interaction with the system (see Section 5.1).

A1.5 QUESTIONS AND POSSIBLE PROBLEMS

A1.5.1 The Changing User Background and Knowledge

How to treat the diversity of user backgrounds, both "statically" (by education and experience) and dynamically (by getting new data and insight)? One of my key points at the San Diego workshop was that "information = surprise", meaning a possibly new dimension in visualisation that might not been adequately addressed so far. It means that the best screen layout or visualisation technique are likely to change very quickly during a session, even for the same set of data.

A1.5.2 The Overwhelming Complexity and Detail

I am afraid that we underestimate the complexity and detail in "user interface theory". The 6th point in 2.1 of Martin's "Guidelines for evaluating ..." is "What possibilities exist for alerting the user?" A meaningful judgment of these other than just stating the types of alarms, demands a lot. Designing a proper alarm system may take the expertise of a complete research institute and more! (An own modest attempt many years ago failed in several aspects!) Section A1.3.6 mentions the numbering of 1419 guidelines. It takes a lot to see the applicability of a "complete set" of HCI guidelines, no matter which set we choose!

A1.5.3 Dilemmas Outside Our Grasp?

My concern here is that "Our Theory" might not be able to tell what is best fronted with specific alternatives. One (for me!) difficult dilemma can be posed from a point made at the San Diego workshop by Ward Page: "good visualisation is more an attention problem than an information transfer problem". Since both 3D and VR tends to "immerse" the viewer, effective information transfer may gain, but attention may loose. What is the right balance? Another dilemma is the use of better symbology where standards are well established. One example here is the NATO symbols describing force units, all being of the same rectangular size (and therefore not easily recognizable at a distance or at a quick glance). User selectable symbology might seem an obvious solution to us. However, I doubt the military community is ready for this! (Are we prepared to suggest and test such a revolutionary principle?!)

A1.6 THE WAY AHEAD?

If the described suggestions are accepted, they seem to at least imply the following tasks before starting evaluating a system:

- Development of checklists (to be started before we meet?)
- Development of "user sequences" (to be started before we meet?)
- Clarification of our "speciality" relative to general "usability" and "usability consultants".
- (To be discussed at our first meeting – or by e-mail?)
- A written procedure to be tried on a system evaluation (formulated at the end of our meeting?)
- Availability of TTCP AGVis description of "visual enhancements" (when?)

A1.7 CONCLUSION

This note has suggested practical measures to bridge a perceived gap between our current theoretical and methodological platform for doing system evaluations; the suggestions to be discussed at our meeting in IST-021/RTG-007 in October 2001. If accepted, the suggestions will imply additional work before we can start evaluating systems; construction of checklists and description of "user sequences" being the more important ones. A clarification of our specialties relative to "ordinary usability consultants" is also recommended. The current non-availability of the candidate system might therefore not be serious at this stage.

Appendix 2 – Update of "Evaluation Procedure" from Aalborg 2001

Tore Smestad, December 13, 2001 (Updated Jan 4 2002)

A2.1 INTRODUCTION

Section 3 of this note contains my update of the "system evaluation procedure" from the meeting in Aalborg, as I promised to do (after the Workshop CfP). It should be seen as a suggestion that might need further adjustments before being applied; we should have a common agreement on a version before we try to test systems with it. Appendix 1 presents a flowchart representation of the procedure as I also promised. The original procedure, as edited by Pavel, is included in Appendix 2 for easy reference.

Important material as a further background are the guidelines for evaluating visualisation systems on www.vistg.net and my note before the Aalborg meeting [Appendix 1 to Annex A, above].

A2.2 ADDITIONS AND CHANGES FROM THE ORIGINAL PROCEDURE

One big addition is to include a "system owner", being a system developer or a person with high knowledge and/or interest in the system. I believe this gives a number of benefits like obtaining more information of the system and its intended use, having a clearer picture of the reader of our "evaluation report", and getting immediate feedback and added information about our system evaluation. This person is involved in the procedure and the reporting (reporting is viewed as part of the procedure).

One serious objection to the original procedure was that the task sequences should not be trained beforehand (point 4.1), since the problems experienced by novices were claimed to be similar to those experienced by expert users during time of stress. Since I still think the level of knowledge and training of the users somehow has to be taken into account, I suggest making a deliberate choice here to allow for different levels – see A2.3 points 1) and 2).

A big change due to the previous point is the "natural creation" of the hierarchy of the task/info/action associated with "Martin's 6 questions". Even if this at the highest level will follow the task sequence described beforehand, the specific way it develops on the possibly many levels below is likely to be quite random. This hierarchy can therefore only be described afterwards, preferably by using a tape recorder for recording intentions, actions, problems, comments and time.

Another change is the use of two passes of the task sequence and associated hierarchy of the task/info/action due to "Martin's 6 questions". In the first pass we simulate a real user, while in the second we behave as ourselves being thoughtful evaluators. Originally, both were done for each "what is the task" item. This division is believed to be more natural and allows for a timing of the first pass (and its steps) which might give valuable information.

A2.3 THE PROCEDURE

Here is the procedure as I suggest it.

1) Prepare the evaluation together with the "owner" by:
   - Decide and agree on the ambition level of the evaluation to be performed.
   - Decide and agree on the type of users and their expertise this evaluation is performed for.
   - Decide and agree on the circumstances of the use of the system in this evaluation
   - (i.e. range/level of stress of users).
   - Obtain or describe 3-15 "tasks"/ "task sequences" to be used in this evaluation.
   - Decide and agree on the documentation level and possibly disclosure of results.
   - Obtain some expected results expressed by the owner.
2) Prepare the evaluation as an evaluator by:
   - Collect and select which HCI/visualisation guidelines to use in this evaluation.
   - Get used to the system to match point 1b, 1c, and 1d (possibly no activity).
   - Prepare the logging of the results in the next steps (checklists, paper logging, tape recording).
3) Go through each of the "tasks"/ "task sequences" individually by the following steps:
   - "Simulate" the completion of the "task"/ "task sequence" by evolving "in a natural way" the hierarchy of the task/info/action associated with "Martin's 6 questions". Do:
     1) Register (tape record?) the answers to "Martin's 6 questions" continually.
     2) Register time information if not automatically done and if decided in 1a.
     3) Make short notes to better recall items for the next point (if not destroying the time info).
   - Recall the created task/info/action hierarchy and register observations in the previous step by:
     1) Try to represent (graph?) the traversed hierarchy of the task/info/actions.
     2) Make a summary (if meaningful) of the types of problems met at points in the hierarchy.
     3) Register violations of specific HCI/visualisation guidelines (see 2a).
4) Check about parallel or interrupted "tasks"/ "task sequences" by:
   - Describe relevant "mixing" of the "tasks"/ "task sequences" (parallel and/or interrupted).
   - Check each of the "mixings" as deemed appropriate, either by the full procedure of step 3 or a subset of it.
5) Register the system features
   - Make a summary of common features of "findings" of points in 4 and in 5.
   - Register lack and present use of "TTCP AGVis visual enhancements" (if selected in 2a).
   - Register the degree of match with "Tore's 13 guidelines" (if selected in 2a).
   - Register good solutions and features experienced.

6) Write the evaluation report relevant to the owner:
   - Organize the chapters as point 1 – 5 of this procedure filled with associated data registered.
   - Make a Chapter 6 suggesting improvements of registered problems if the improvements can be motivated by violated guidelines; collect/group the rest of the observed problems.
   - Make a brief overall summary including comments to point 1f.
7) Write "evaluation-evaluation" observations and conclusions for this evaluation (as a new chapter or as a separate report; data from step 6b could be of special interest as is a grouping of the types of observed problems and to which extent "Our Theory" is involved).

A2.4 EVALUATION REPORTS

Each evaluation should produce some form of a written evaluation report to the system owner. The inclusion of the owners expected results may give two benefits: the owner may get more involved and the results may be more "pure information" to him (remember: "information" is defined related to expected outcome.) One chapter of the report, or a separate report, should be for ourselves describing lessons learned and experiences about the procedure and methods we used.

The content should be organized according to the procedure, see more suggestions in my note of Aug. 23 and point 6 of the procedure. I suggest to only present possible solutions to experienced problems as long as they might be supported by some guidelines. The other problems will probably constitute an incitement to produce new guidelines.

A2.5 SOME COMMENTS

The coming evaluations are, at least to begin with, aimed more for our own benefit rather than for the owners. However, I expect the presence of the owners and the task of making evaluation reports to them, add realism which will be quite valuable for us too.

I still worry about what our specialties should be, now I have a creeping feeling that we are doing things that most people could do without any interest or knowledge in visualisation as we define it. It looks to me that we are dealing with the general usability of systems. Maybe one could view general usability as the lower part of a pyramid, like a "Maslov hierarchy" with "visualisation specialties" on the very top. Hopefully we can get to treat those questions as well!

I think to remember a philosopher claiming (Heidegger?) that we observe only the imperfect aspects of the world. That may have a bearing on point 3b1 in that the lower levels of the hierarchy being registered (each branch initiated by "what is the task") are mainly the ones causing problems.

The "naturally created" hierarchy is expected to be random in the sense that if an other similar person did it, or the same person did it at an other time or once more, the hierarchy would change. To investigate the latter might tell something about learning; this is not included in the procedure, but might easily be done.

I think using tape recording is beneficial, almost a must, to be able to register the hierarchy and its timing without a disturbing influence from the registering process. Video could also be useful, but might be an "overshoot" at this point.

A2.6 CONCLUSION

This note has suggested an updated evaluation procedure relative to the one we discussed in Aalborg. Several changes are made and new aspects are introduced; one important one is the inclusion of a “system owner”. The procedure is meant as a common framework for doing some initial “visualisation evaluations” early in 2002 on selected systems. The procedure presented in Chapter 3 and represented in Appendix 1 more visually, should be critically inspected. We should try to agree on a common procedure for our initial evaluations.

Annex B – VISUALISATION RESEARCH AND DEVELOPMENT IN THE UNITED STATES

S. Chipman
U.S. Office of Naval Research

Attempting to survey visualisation research in the United States is a daunting task. Realistically, it is not a task that can be done without significant resources, and these were not available. I searched several readily available databases to provide an overview of work that is going on in the United States.

B.1 DEFENSE TECHNICAL INFORMATION CENTER DATABASE

As an initial attempt, I searched the Defense Technical Information Center file of unclassified scientific and technical reports. A search on the term *visualisation* resulted in 250 hits, a substantial fraction of which I judged to be relevant when I examined them. A large fraction of these reports are about battlespace *visualis*ation involving computer displays, although there also seems to be an older usage in military training that refers to the commander's mental capacity to *visualis*e battle situations. Many of the reports are concerned with comparing the value of 2D and 3D displays of various kinds. Although his project was not among those that surfaced in the DTIC search, Dr. Greg Trafton of the Naval Research Laboratory in Washington, DC has been engaged in a theoretical effort to systematize the disparate results of such studies, to determine what factors lead to empirical superiority of 2D or 3D displays. Quite a few other projects are concerned with the display of inherently 2D or 3D information, such as the display of terrain or weather information. Isolated projects were concerned with such issues as the display of network security status or representing uncertainty. Several reports represent an ONR-funded program of research by a mathematician, Edward Wegman of George Mason University, who has been concerned with *visualis*ation methods for exploring high dimensional data.

One single report appears to contain more information than the rest put together. ADA401261, Advanced Displays and Interactive Displays Report Compendium III, 2002, "contains more than 300 citations and abstracts of papers and presentations produced by the Advanced Displays and Interactive Displays consortium during the 5-year U.S. Army Federated Laboratory program. … The … consortium seeks to provide innovative, cost-effective solutions to information access, understanding, and management for the soldier of the future. The research encompasses a range of topics. Some work concerns the representation of uncertainty and imprecision in databases or the representation of relationships in multimedia databases, in ways that are compatible with human cognitive-processing capabilities. Other work adopts the means of human communication (such as speech, gesture, eye gaze and lipreading) for human-computer interaction. Additional work explores methods for incorporating information in virtual reality displays that support decision-making without distracting or overwhelming the soldier. Although diverse, the research is linked by its overriding goal: the presentation of information in a form that allows effective human understanding and decision making in complex battlefield situations." {abstract}

It is worth noting that DTIC is becoming less useful as an information resource even as it has become much more accessible. DARPA has not submitted project information for years. ONR has ceased to submit information about new projects. Thus, there may be many more military supported projects in *visualis*ation, even unclassified ones, that do not show up in the data base. On the other hand, some reports from foreign sources do appear in the database. ADA360645 is an Australian effort that may be of interest: Neelam Naikar, Perspective Displays: A Review of Human Factors Issues, Feb 1998, Defence Science and Technology Organisation, Canberra, Australia. Another Australian report also seems to be of potential interest. ADA355765 by Michael D. Lee and Douglas Vickers, Psychological Approaches to Data Visualisation, Adelaide University, July 1998. The abstract of this report starts with a nice statement of purpose: "The aim of data *visualis*ation is to display a body of information in a way which allows accurate and effortless human comprehension and analysis."

B.2 IEEE VISUALISATION CONFERENCES AND OTHER SIMILAR CONFERENCES

These conferences are probably the best up-to-date source of information about visualisation research in the United States, although, of course, they do have international participation. **VIS2003** took place October 19-24, 2003, in Seattle Washington. It was the 14th annual conference – see http://vis.computer.org/vis2003/. Programs for most of these conferences are available on the web and conference proceedings are published. There are now two more specialized conferences associated with this conference: InfoVis, the IEEE Symposium on Information Visualisation and PVG, the IEEE Symposium on Parallel and Large-Data Visualisation and Graphics. This is the ninth year for the former and the sixth year for the latter. Programs for these are also available on the web. There is also a related IEEE journal, *Transactions on Visualisation and Computer Graphics*. This journal is now in its ninth volume. There is also an IEEE Technical Committee on *Visualis*ation and Graphics. Their site says they welcome members who want to stay in touch with these fields from around the world. The current url is http://www.cc.gatech.edu/gvu/tccg/. Psychological research on the use or usability of visualisations is relatively rare at these conferences, but does occur.

Recently, a supplementary issue of the Proceedings of the U.S. National Academy of Sciences that is largely devoted to the visualisation of information has been published [1].

There are also several other series of visualisation conferences. There is a series of conferences called Visualisation and Data Analysis affiliated with SPIE (International Society for Optical Engineering) and IS&T (The Society for Imaging Science and Technology). The next conference is scheduled to take place in San Jose, CA, January 17-20, 2005. The call for papers states that the conference covers all aspects of data analysis and visualisation as well as issues affecting successful visualisations. The conference has grown rapidly over the years and has attracted participants from throughout the world. According to one of the conference chairs, the SPIE conference generally has less evolved results than the IEEE series. There is also a series, The International Conference on Visualisation, Imaging and Image Processing (VIIP), sponsored by IASTED (International Association of Science and Technology for Economic Development). The same organization also sponsors a series of conferences on Computer Graphics and Imaging. CGIM 2004 will take place in Kauai, Hawaii, August 17-19, 2004.

Discussions of visualisation may also occur at the supercomputing conferences, such as SC2001, where there was a plenary session entitled, "Scientific visualisation: Bridging the complexity threshold," by Chris Johnson of the University of Utah. I presume SC stands for supercomputing, although the web pages characterize this as "The International Conference for High Performance Computing and Communications" and the "SCxy conference series."

B.3 PSYCINFO

PsycINFO is the database of psychological research literature maintained by the American Psychological Association. It is quite comprehensive, including a very wide range of psychological journals, even some that are published in languages other than English, as well as books, dissertations, and some technical reports. For example, technical reports that are in DTIC may also appear in PsycINFO. Obviously, the publications in PsycINFO are not restricted to authors from the U.S. A search of PsycINFO revealed that not a great deal of research related to visualisation is being reported in the psychological research literature. Much of what is there is about educating people to use graphs and charts or about their use in statistical analyses by psychologists. One U.S. effort of likely interest to the RTG did show up. E. Morse and M. Lewis of the National Institute of Standards and Technology (NIST) published a paper, "Evaluating visualisations: Using a taxonomic guide." International Journal of Human-Computer Studies, 2000, 53, 637-662. According to the abstract, this was an effort to test visualisations with a general visual taxonomy. A goal was to avoid underestimating the capabilities of newer systems by evaluating with tasks that can be accomplished by older, less capable systems as well.

B.4 U.S. NATIONAL SCIENCE FOUNDATION (NSF)

Visualisation appears to be a relatively large area of investment for NSF. Of course, NSF is the agency supporting supercomputer centers in the U.S., which implies substantial involvement with visualisation technology. In 1987, NSF published a report (McCormick, DeFanti, and Brown, Visualisation in Scientific Computing) that promoted the importance of research on visualisation, although this did not seem to result in the desired specific program on the topic. More recently, there is an NSF report on Terascale and Petascale Computing that is available on the NSF web site (nsf.gov). This has a major section, "Data Visualisation Challenges and Opportunities", that argues for even larger investments in visualisation research, especially directed at very large datasets. This sums up a series of 3 workshops that were sponsored jointly by the Department of Energy and NSF in 1998.

An on-line database of the abstracts of awards made since 1989 is available on the web. A search on the term *visualisation* yielded 2866 hits. These awards are quite diverse. Among the most interesting to the RTG might be these current awards:

> Abstract #0086065: Visualisation of Multi-Valued Scientific Data: Applying Ideas from Art and Perceptual Psychology. This is a $2.3M grant to David Laidlaw, Michael Tarr, and George Karniadakis of Brown University.
>
> Abstract # 0222991 Visualisation: A metadata Driven Interface Technology for Scientific Data Exploration. A $265K grant to Kwan-Liu Ma and Michael Gertz of the University of California at Davis. This abstract says the project is trying to facilitate effective reuse, sharing, and cross-exploration of visualisation information.

There are a number of CAREER grants to relatively young investigators that combine research on visualisation with the development of courses or entire curricula to train students in visualisation. An interesting example of these, which seems to have been renewed is

> Abstract # 0238261: Visualising Knowledge Domains. This is a $400K grant to Katy Borner of Indiana University.

Another of these is

> Abstract #0092308: Assisted Navigation in Large Visualisation Spaces to Christopher Healey of North Carolina State University.

There are a large number of grants for instrumentation of visualisation laboratories, either for research or education purposes. There are also grants from the education section of NSF for the use of visualisation in teaching various kinds of scientific subject matter, such as physics, weather science, and experimental psychology. Of course there are also research grants for visualisation in particular sciences. There was a grant for a two week training workshop in visualisation conducted at the Illinois super computing center.

A further search on *visualisation AND psychology* yielded only 30 hits, but was not more informative. The Brown University grant mentioned above topped the list. There were grants for visualisation of psychological data, including fMRI imaging, and educational grants. Notably there really were not grants for psychological research on visualisation like that being done by Greg Trafton at the Naval Research Laboratory.

B.5 U.S. DEPARTMENT OF ENERGY

As noted above, the Department of Energy co-sponsored several recent workshops on visualisation with NSF. The Department of Energy has responsibility for nuclear weapons and nuclear power in the United States. Consequently, its numerous laboratories have many supercomputers. The very first Cray ever built, for example, was installed at the Los Alamos laboratory. As the nuclear line of business has declined, the laboratories (e.g. Los Alamos, Oak Ridge, Sandia) have sought new missions and have become involved in a wide variety of activities. Visualisation R&D in the Department of Energy is in the Computer Science program, itself part of the Mathematical, Information, and Computational Sciences, which is described as a subprogram of the Advanced Scientific Computing Program. Here is the description of the Computer Science Program: This includes research in computer science to enable large scientific applications through advances in massively parallel computing, such as very lightweight operating systems for parallel computers, distributed computing such as development of the Parallel Virtual Machine (PVM) software package that has become an industry standard, and large scale data management and ***visualisation***. The development of new computer and computational science techniques will allow scientists to use the most advanced computers without being overwhelmed by the complexity of rewriting their codes every 18 months.

There is a National Energy Research Scientific Computing Center (NERSC) at the Lawrence Berkeley National Laboratory, which is characterized as one of the nation's most powerful unclassified computing resources. This center provides both computing resources and intellectual services such as the development of innovative algorithms, simulations and ***visualisation techniques.*** Access to NERSC is available from anywhere in the world through ESnet, which provides very high bandwidth service. NERSC has a Visualisation Group, a dedicated visualisation server, and uses the NCAR (National Center for Atmospheric Research) graphics library. Visualisations can be performed remotely. MPEG movies of visualisations can be made. A key member of the staff in this area seems to be Ravi Malladi (malladi@euphrates.lbl.gov). NERSC is now physically located at the Oakland Scientific Facility in downtown Oakland, California.

NERSC issues calls for proposals in the INCITE (Innovative and Novel Computational Impact on Theory and Experiment) program, in which outside investigators can request large allocations of computing resources. Eligible investigators are those engaged in scientific research with the intent to publish results in the open peer-reviewed literature. The program specifically encourages proposals from universities and other research institutions without any requirement of current sponsorship by the Office of Science of the Department of Energy (sponsor of NERSC).

All of the above information, except for the first few introductory sentences, comes from publicly available web sites, such as www.science.doe.gov/ascr or www.nersc.gov.

B.6 NASA

The primary locus of visualisation R&D at NASA appears to be the Numerical Simulation Facility (NAS) at the NASA Ames Research Center in the San Francisco Bay Area. The Visualisation Lab staff consult with and assist researchers in creating animations, videos, and acoustical data to help analyze and present scientific results. Their visualisation products focus on meeting the data analysis need of aerosciences researchers, but actual examples seem to be somewhat more diverse than that suggests. For example, they have developed a prototype application called "growler" (formerly Virtual Mechanosynthesis) that provides users with a tool to explore molecules and materials they are working with. The prototype growler environment includes COSMOS (Computer Simulations of Molecular Systems), a scientific code used to stimulate chemical and biomolecular processes such as protein folding, protein-DNA interactions, and functions of cell membranes.

Among the interesting projects found on their web site (www.nas.nasa.gov – but it warned that the location would change) were these:

> Kristina D. Micelli (1994), A Data-Centered Framework for an Assistant-Based Scientific Visualisation System. NAS Technical Report NAS-94-006, December, 1994.
>
> "This research presents the development of an assistant-based visualisation system, based on a data-centered framework that supports the rapid production of effective graphics for scientific visualisation. The framework is composed of a data model, a user model and a machine model that represent the scientific user and the software/hardware environment. Each model contains knowledge related to the production of effective visualisations, including rules from perception, graphic design, and the scientific domains. … a prototype assistant called MDV, or MultiDisciplinary Visualiser, was developed … The application domain was computational aerosciences, involving multidisciplinary data from computational field dynamics and structural dynamics. The system was developed with the close involvement of scientists performing the simulations." This seems to have been a dissertation research project.

MARSOWEB provides images of the planet MARS in an interactive data map viewer developed by a researcher named Deardorff. There is also a VRML component which provides a 3D image of Mars' surface. Users can zoom through Mars' canyons and valleys or fly over its volcanoes and desert dune fields. The background images are from Viking, overlaid by higher resolution images from the Mars Orbital Camera on board the Mars Global Surveyor. Users can also interact with data from the Mars Orbiter Laser Altimeter, an elevation measurement tool on board the Global Surveyor. The original purpose was to assist in selection of landing sites on Mars, but the site continues to be enhanced with additional data.

Hyperwall is a 7 x 7 assembly of 49 flat panel screens, each driven by its own dual-processor computer with high end graphics card, to provide 49 different viewpoints on complex, high dimensional datasets. Rather surprisingly, the claim is made that the choice of the number of screens was based partially on George Miller's famous paper, "The magical number seven, plus or minus two: some limits on our capacity for processing information." The description even provides a link to a copy of that paper!

The greatest concern for human usability of visualisation seems to be expressed by an individual named Al Globus, who produced a tech report that is a short position paper on the need to test visualisation software more rigorously to ensure that it works and to conduct experiments with human subjects to evaluate visualisation techniques. He produced another technical report in 1994 entitled, "Principles of Information Display for Visualisation Practitioners" that is a summary of Edward Tufte's books on the subject. It does not appear that serious research on the usability of visualisation systems has been a feature of the NAS Ames program.

B.7 OFFICE OF NAVAL RESEARCH AND NAVAL RESEARCH LABORATORY

The Naval Research Laboratory in Washington, DC, has a supercomputing facility. Like other supercomputer centers, scientific visualisation is a prominent activity there. Larry Rosenblum, an applied mathematician on the NRL staff, has been a leading figure in the developing scientific visualisation community. Fluid dynamics has been an important application area for NRL, not surprisingly. Like NSF, but on a much smaller scale, the computer science and mathematics programs at ONR have sometimes supported visualisation work. Wegman's research program, mentioned in the DTIC section above is an example. At one point, I proposed an interdisciplinary program together with Rosenblum and a program officer in the mathematics area, that would have emphasized psychological research related to scientific visualisation, attempting to understand how visualisations are perceived and used, what features attract the user's attention as being interesting. An eventual goal would have been an automated system that would extend the user's perceptual capabilities and augment them with automated exploration of large dimensional data spaces such as those Wegman has been interested in. This proposed program was largely inspired by the lack of attention to psychological research and theory in the 1987 NSF report, which made vague reference to the somewhat mystical and unanalyzed powers of the human eye. It was also heavily influenced by the fact that selected illustrations of visualisations almost always seem to involve strong symmetries or minor departures from symmetry. However this program was not selected for funding so there has been no substantial investment in the area. Greg Trafton (work described below) has been supported both by internal NRL funds and by funds from a human computer interaction program at ONR. Other reports are in [2,3].

B.8 SOME INDIVIDUAL RESEARCH PROJECTS/PROGRAMS OF NOTE+

B.8.1 Multi-Modal Interfaces for the DARPA Command Post of the Future

Ken Forbus and Brian Dennis of Northwestern University. The project description on their web site says, "We will create a sketch understanding system that can be used in multimodal interfaces for constructing and evaluating courses of action via map-based sketching and speech. We will implement analogical processing systems for supporting sketch understanding, elaboration of courses of action based on libraries of operations plans, visualisation of similarities and differences involving alternate COAs and the evolving situation versus plans in execution, and for suggesting higher-level explanations for enemy actions." At least the sketch understanding system exists at this point; I have seen it demonstrated. Not surprisingly, military commanders find it very appealing – natural and simple to use.

B.8.2 Interactive Multi-Sensor Analysis Trainer (IMAT)

This applied research program has now been going on for many years and has evolved from being a training project to also being a decision-aiding system, or family of decision aiding systems. IMAT is an outstanding application of scientific visualisation techniques to assist people in understanding the complex interactions among sound energy, the ocean environment, and the sensor systems available to detect that sound energy. These interactions are indeed complex, too complex to be mentally simulated even by those who understand them very well. And the details matter in terms of being able to detect an enemy or having one's own ship be detected. Examples of the graphics used in this family of systems are currently available on the web site describing ONR's Training-Related R&D – see http://www.onr.navy.mil/sci_tech/personnel/cnb_sci/342/majapps/majapp.htm. These systems are now in wide use for all kinds of sonar training. The one evaluation study that was done had spectacularly positive results. Previously sonar training consisted largely of memorizing a great many facts without much organization or meaning, so the opportunity for improvement was great. An unclassified paper describing this work appears in a NATO Symposium on Advanced Training Technologies [4].

B.8.3 Knowledge Wall/Knowledge Web/Knowledge Desk

For a number of years, SPAWAR Systems Center San Diego, supported by Pacific Science and Engineering Group, Inc., has been working on this effort to support shared situation awareness and decision-making in the Joint Operation Center. A wall-sized shared display – or "Knowledge Wall" – fusing all information relevant to mission status, has been proposed as a solution to this need.

Somewhat similar efforts are examining the information requirements for decision-makers in other military command centers. The Air Force is currently developing a multimedia "Interactive Datawall" to provide fused, multi-source information for C4I environments, and SSC San Diego is also developing a Tactical Situation Awareness Tool (TacSAT) for the U.S. Marine Corps. The design of a knowledge wall must take into account the specific needs of its intended users, so there are important differences as well as similarities among these efforts.

In order to determine what features and content should be in this new display, structured interviews were conducted with JOC command elements . Most of the participants were senior level personnel, admirals and battle watch captains, but individuals producing information were also included. Fourteen requirements were distilled from the interviews with potential knowledge wall users. The Knowledge Wall, they said, must support:

- Shared situation awareness among its JOC users;
- The integration of relevant mission status information;
- An intuitive graphical interface;
- Consistently formatted information;
- A tactical focus for the displayed information;
- The display of information to supplement tactical data;
- The display of mission goals and Commander's Critical Information Requirements;
- The display of summary information provided by "anchor desk" or support staff;
- The ability to connect and coordinate or collaborate with others at diverse locations;
- Support tools and aids for challenging cognitive tasks;
- A flexible configuration that can easily be changed by users;
- The ability to drill down through displayed information for more detail;
- Display of information age and reliability; and
- Tactical overlays to highlight different types of information.

Many responses referred to the need to support difficult cognitive processes by such aids as directing users' attention to important information and alerting them to important changes in information; highlighting important information and "decluttering" irrelevant display information; supporting management of planned actions; and highlighting potential conflicts between domains or ambiguous and conflicting information.

In many respects, such as the use of tactical overlays, the knowledge wall design is building upon and integrating many past R&D efforts in display design.

The Knowledge Wall (Figure B-1) features a series of windows incorporating decision support tools tailored to the CJTF, as well as windows with "summary status" information being "pushed" by the liaison officers representing the various CJTF departments such as weather, air, intelligence, etc. The battle watch captain can choose aspects of the situation on which to focus, and the appropriate displays for the collaborative workspace in the center of the data wall.

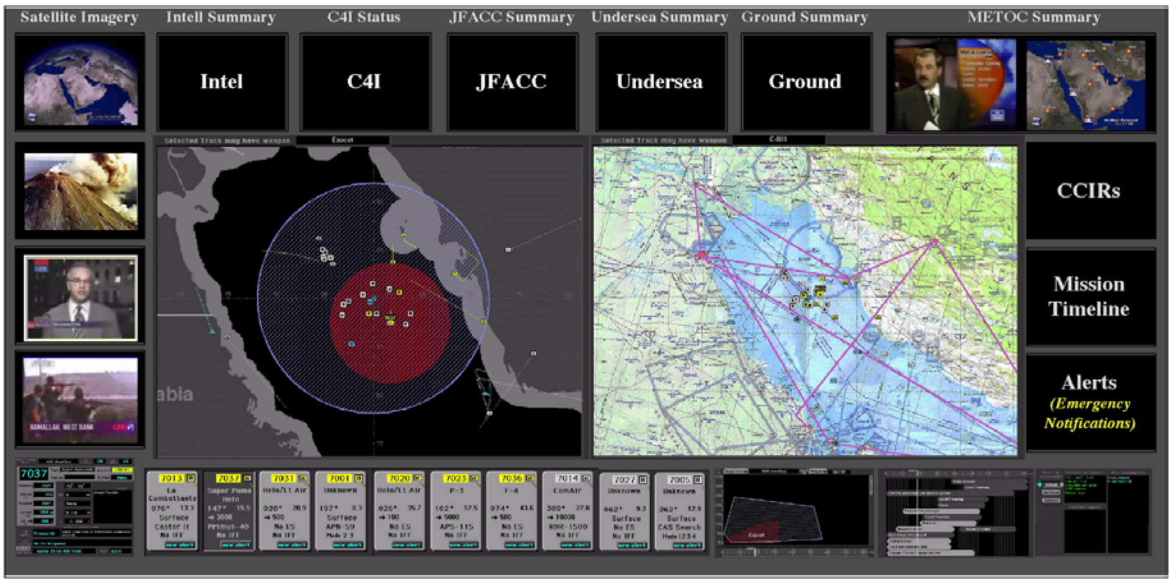

Figure B-1: Knowledge Wall.

The Knowledge Web complement to the Knowledge Wall is that "knowledge" created by the command staff can be captured in and distributed in real time using Web pages, in place of the traditional daily and/or watch-turnover briefings typical today. This change improves the speed of command. Template-based authoring tools have been developed to support the development of summary pages for each type of information. It is important to understand that the "knowledge wall" is not just a collection of displays, but also an interface to a human organization in which different individuals have responsibility for processing different kinds of information needed for command decision-making. Newer efforts that appear similar, such as an SAIC effort for the Department of Homeland Security, will need to develop an appropriate background human organization as well as the computational and display facilities.

The large Knowledge Wall display is supplemented by Knowledge Desks (Figure B-2) which also incorporate multiple displays (4-6) designed for use by individuals or smaller groups. These provide a content production workspace, but also access to overall status displays, providing important context to those who are producing summaries for their area of responsibility. They also provide a way of dealing with the issue of size and visibility of information on the possibly distant large display.

Figure B-2: Knowledge Wall with Knowledge Desks.

Data collected from Global (war game) 2000, an operational exercise environment, demonstrated that the Knowledge Web could dramatically increase situational awareness and the speed of command. After this exercise, some modifications were made and the first operational implementation of the Knowledge Web occurred onboard the USS Carl Vinson in May 2001, for use during their WESTPAC deployment. Later, the Carl Vinson Battle Group became the Composite War Commander for military activities supporting Operation Enduring Freedom (Afghanistan). This provided the opportunity to assess how the Knowledge Web supports users' requirements within critical operational environments. Data were collected using multiple methods – automated records of use, interviews and surveys.

For at least some users, the value of this system became strikingly evident under these pressured circumstances, as it had not been before. Commanders felt they had access to current information at all times and felt that traditional briefing sessions became more interactive, collaborative problem solving sessions.

Not surprisingly, a destroyer squadron that accessed the system information, but without any training provided or on-line help, and who were also hindered by unrelated technical problems, found the system considerably less useful. In particular, they did not understand the symbols and colors that were used to indicate changed, critical information.

The investigation of practical usability issues continues. For example, the templates for authoring summary pages were found to be easy to use, but more flexibility in format was desired. On the other hand, consumers of the information found uniformity of format to be very helpful in accessing and recalling needed information.

Project reports are [5 – 7].

POC for the project: Dr. Jeffrey G. Morrison, Space and Naval Warfare Systems Center, San Diego, 619-553-9070, jmorriso@spawar.navy.mil

B.8.4 Eye-Tracking Methods for Evaluating Complex Displays

This research has been carried out by Prof. Sandra Marshall of San Diego State University, a cognitive psychologist and expert in eye-tracking research, in collaboration with Dr. Jeffrey G. Morrison of SPAWAR. Marshall has tracked the eye movements of experienced military personnel while they perform in simulated scenarios using experimental CIC displays. The experimental set-up and the eye-tracking results are depicted in Figure B-3 below. The results certainly demonstrate that the information is not being taken in "at a single glance." They are being related to a hybrid schema theory of anti-air warfare tactical decision making that Marshall developed earlier.

Figure B-3: Illustration of Eye-Tracking Data Showing the Difference between an Inexperienced (yellow) and Highly Experienced (red) Watchstander.

B.8.5 Force TEWA

For many years, the Johns Hopkins Applied Physics Laboratory, a laboratory funded by the U.S. Navy, has been working on the development of technically advanced displays intended to improve situation awareness in anti-air warfare. The relevant program was originally called the Battle Group AAW Coordination program. Its best known product is probably Force Threat Evaluation and Weapons Assignment (Force TEWA) prototype, featuring realistic icons in 3D displays (Figure B-4). This prototype has evolved into the Area Air Defense Commander (AADC) prototype. This is intended to be capable of integrating Navy, Army, Air Force, and Marine Corps air defense elements to produce a coordinated AAW station plan and support the AADC in real-time tactical operations. In addition to the laboratory prototype, a second copy was installed in an Aegis cruiser for demonstration and testing in an at-sea exercise. The AADC was very well received and is going into production for the Navy.

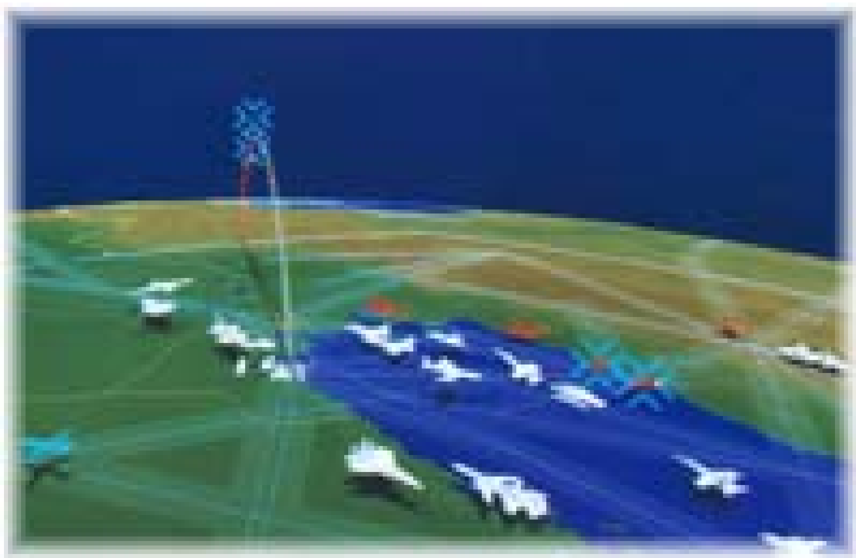

Figure B-4: Example of 3D Situation Display Featuring Realistic Icons.

B.8.6 TOPSCENE

TOPSCENE (Tactical Operational Scene) is an operational training (that is, mission rehearsal) system consisting of 3D imagery products and infrastructure (Figure B-5). It was produced for the Naval Air Systems Command by Lockheed Martin Missiles and Fire Control, and is intended for use by the U.S. Navy, Air Force, and Army. LM corporate advertising on U.S. television has shown imagery from this system. It is claimed to provide rapid, accurate database construction and real-time 3D fly through capability. The databases are acquired from various sources (photographic, sensor, satellite, etc.) and processes by the off-line TOPSCENE Data Base Generation Systems. Two dimensional imagery is converted into three dimensional “fly through” and “walk through” battlefield visualisation simulation scenarios. A data base can be created from images in about two hours. They are stored on removable digital disks and transported to the user. The operator station is menu-driven and controlled through the key board. Mission rehearsal is performed using a throttle and control stick; other weapons systems controls can be added. The operator can select fixed wing aircraft motion, helicopter motion, guided missile seeker motion, or walking. The Air Force version of TOPSCENE also permits specific training for crews using guided munitions. It provides bull’s-eye scoring, miss distance accuracy reports and a best score notification during training. System enhancements permit changes in visibility, time of day, sensor imagery, graphics overlays, instrument indications and even threat warnings. TOPSCENE 4000 is intended to be able to replicate the precise visuals encountered in any environment.

Figure B-5: TOPSCENE Operational Training System.

Contacts:

- Lockheed Martin Missiles and Fire Control, Business Development, 407-356-4464; and
- Naval Air Systems Command, Patuxent River, Maryland, 301-757-8136.

B.8.7 SPAWAR/Pacific Science and Engineering Group, Inc.: 2D vs. 3D Displays

For a number of years, researchers at the Space and Naval Warfare System Center in San Diego and at the Pacific Science and Engineering Group, Inc., have been researching issues related to three-dimensional displays and icons, such as the Force TEWA display described above, as contrasted to more conventional two dimensional tactical displays with abstract icons. Although the 3D displays are very appealing to the customer community, it appears that task performance is most often superior when based on 2D displays. In a recent oral report on the research, Smallman of this research group conceded that 3D displays are so popular that they seem unstoppable and suggested a change of research focus to amelioration of the perceptual problems with 3D displays.

Reports are at [8,9].

B.8.8 PARC Research on Information Visualisation

For a considerable period of time, Stuart Card and others at the Xerox Palo Alto Research Center, now just PARC, an independent organization, have been working on finding ways to use the new graphic capabilities of personal computers to lower the cost of finding information and accessing it once found. They believe that the structure of information, 3D technologies, interactive animation, and the human perceptual system can be effectively exploited to improve management of and access to large information spaces. Cognitive processing load can be shifted, they believe, to the perceptual system. These statements are paraphrased from one of their publications, Robertson, Card and Mackinlay (1993) Information Visualisation Using 3D Interactive Animation, Communications of the ACM,36, 57-71. Card is one of the primary figures behind the InfoVis Symposia that are now attached to the IEEE Visualisation Conferences. Currently Card and his associates are focusing on what they call information visualisation tasks. For example, they have studied how the publisher of a computer industry insider newsletter does his work. Now they are working with intelligence analysts to see how they perform their tasks and how their work can be facilitated with computer tools.

B.8.9 Trafton's Research at NRL

Greg Trafton made a presentation to the RTG when we visited NRL. He talked about a detailed observational study of Navy weather forecasters preparing their daily weather brief, which involves the use of several graphic resources of weather information. The entire task, which involves two individuals of varying experience level, was video-taped and analyzed in extreme detail, tracing every item of information in the final brief to its origin in the information sources. One interesting finding was that the forecasters seemed to prefer prediction of future weather on the basis of their own mental models and did not use animated models available to them [10]. Trafton has also done similar work with scientists using scientific visualisations. In addition, he has been engaged in a theoretical effort to explain varying results concerning the superiority of 2D or 3D displays.

Trafton's papers are available on http://www.aic.nrl.navy.mil/~trafton/visualisation.html

B.8.10 NOAA Science on a Sphere

In the July 7, 2003 issue of *Federal Computer Week*, there is an article describing this new system. SOS uses computers and video projectors to display animated geophysical and atmospheric information on a large, fiberglass globe. It has sophisticated software that can convert a wide range of scientific information into a format that the globe can display, singly or in combination. It can also incorporate other kinds of data such as global satellite imagery. Unlike other kinds of displays, the mapping of information is not distorted. Even experienced meteorologists claim to have gotten new insights from the undistorted display, such as the relationships between wet and dry zones in the tropics. At present, the system is intended primarily for educational museum use, but the military has also shown some interest in it.

B.9 REFERENCES

[1] Shiffrin, R.M. and Börner, K. Mapping knowledge domains, PNAS 2004 101: 5183-5185; published online before print January 23, 2004.

[2] Chipman, S.F. Towards a Synthetic Eye: Psychological Issues in Data Visualization. Unpublished manuscript, November 10, 1992.

[3] Wegman, E.J. and DePriest, D.J. (Eds.), (1986) *Statistical Image Processing and Graphics*. New York & Basel: Marcel Dekker, Inc.

[4] Wulfeck, W.H., Wetzel-Smith, S.K. and Dickieson, J.L. (2004) Interactive Multisensor Analysis Training. Paper presented at the NATO RTO Human Factors and Medicine Panel (HFM) Symposium on

Advanced Technologies for Military Training, 13-15 May 2003, Genoa, ITALY. Neuilly-sur-Seine, Cedex, France: North Atlantic Treaty Organisation Research and Technology Agency (ISBN: 92-837-1117-3).

[5] Smallman, H.S., Oonk, H.M., Moore, R.A. and Morrison, J.G. "The Knowledge Wall for the Global 2000 War Game: Design Solutions to Match JOC User Requirements".

[6] SPAWAR Systems Center San Diego: Technical Report 1860, August 2001. Approved for public release; distribution unlimited.

[7] Rogers, J.H., Oonk, H.M., Moore, R.A., Averett, M.G. and Morrison, J.G. "The design, implementation and use of Web-technologies to facilitate knowledge sharing: A "real-world" application".

[8] Smallman, H.S., Schiller, E. and Mitchell, C.A. (2000) Promoting Rapid Situation Awareness in Tactical Displays: The role of 3-D perspective views and realistic symbols. *Proceedings of the IEA 2000/HFES 2000 Congress*, pp. 3-394 – 3-397.

[9] Smallman, H.S., St. John, M., Oonk, H.M. and Cowen, M.B. (2000) When beauty is only skin deep: 3-D realistic icons are harder to identify than conventional 2-D military symbols. *Proceedings of the IEA 2000/HFES 2000 Congress,* pp. 3-480 – 3-483.

[10] Trafton, J.G., Kirschenbaum, S.S., Tsui, T.L., Miyamoto, R.T., Ballas, J.A. and Raymond, P.D. Turning pictures into numbers: Extracting and generating information from complex visualizations. International Journal of Human-Computer Studies. 2000 November Vol. 53(5) 827-850.

Annex C – GERMAN WORK ON VISUALISATION

A. Kaster
FGAN-FKIE

C.1 SURVEY OF GERMAN VISUALISATION ACTIVITIES (2003 – 2004)

In Germany, research on Visualisation of massive datasets is spread across a diversity of civil and military agencies and institutions. Most military projects are coordinated by the German Ministry of Defense and subordinate agencies. They are carried out by different research institutions, universities and defense companies.

The following summary presents a selection of visualisation activities carried out in Germany in 2003 – 2004 including projects/studies/research. The main focus is put on the visualisation activities in the Research Institute for Communication, Information Processing and Ergonomics (FGAN-FKIE) as well as several academic visualisation activities.

References and descriptions on activities are generally structured into performing institution, reference of institution, name/title of visualisation work, reference and type of visualisation work as well as a description.

Institution:

C.1.1 Research Institute for Communication, Information Processing and Ergonomics (FGAN-FKIE)

Providing scientific and technical expertise for the German MoD and subordinate agencies in the area of C4ISR and innovative Human-System-Interaction – www.fgan.de

C.1.1.1 Knowledge-Based Human-Machine Interface for Future Naval Combat Direction Systems [1]

Future command and control systems for naval ships must particularly meet the specific requirements emerged from the shift from blue-water scenarios of the past towards multinational peacekeeping missions, joined and combined operations and littoral warfare. Rapidly changing situations found especially in the Anti-Air warfare require decisions to be made in high dynamic and complex, i.e. mixed, environments. Since a support of operators by human personnel is out of discussion for financial reasons and is associated with questionable prospects of success, an essential basis for operator performance and situation awareness is an ergonomically optimized and operator-adaptive human-machine interface that facilitates all aspects of human handling. Supporting the cognitive phase of decision-making by means of computer-based assistance in the form of decision-support systems is another important aspect of future combat direction systems. A concept for an enhanced knowledge-based, human-machine interface for future combat direction systems of naval ships has been developed, including the underlying models, techniques, and presumptions. Theory to practice can be given.

Example: Representation of decision relevant information of tracks in the airspace for situation assessment and evaluation. Integration of various information by use of polar diagrams (Figure C-1).

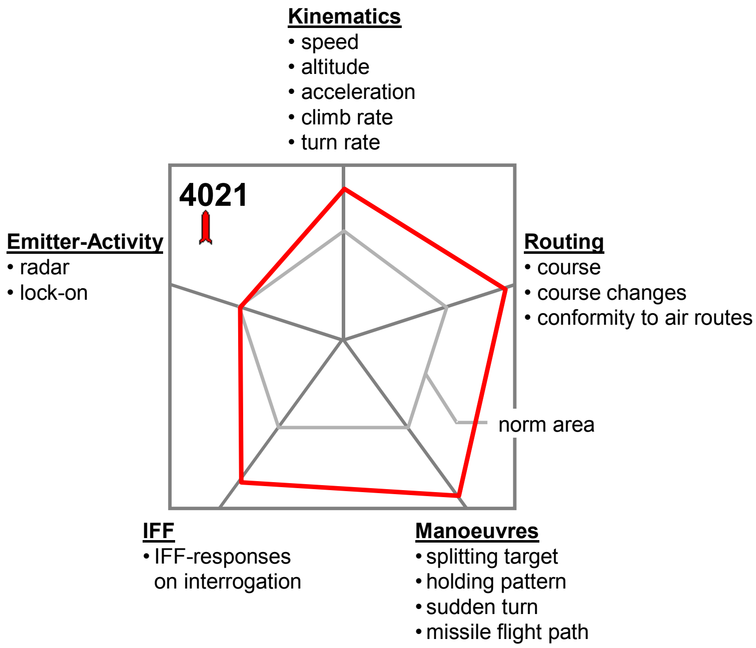

Figure C-1: Draft of a Polar Display for Visualisation of Air Track Attributes.

The Tactical Situation Display positioned in the central visual field of the operator provides essential information necessary for situation evaluation and handling (Figure C-2). It shows all sensor-detected dynamic objects (DTOs) overlaid to a chart of geo-referenced information which is provided by an ENC embedded in the computer network.

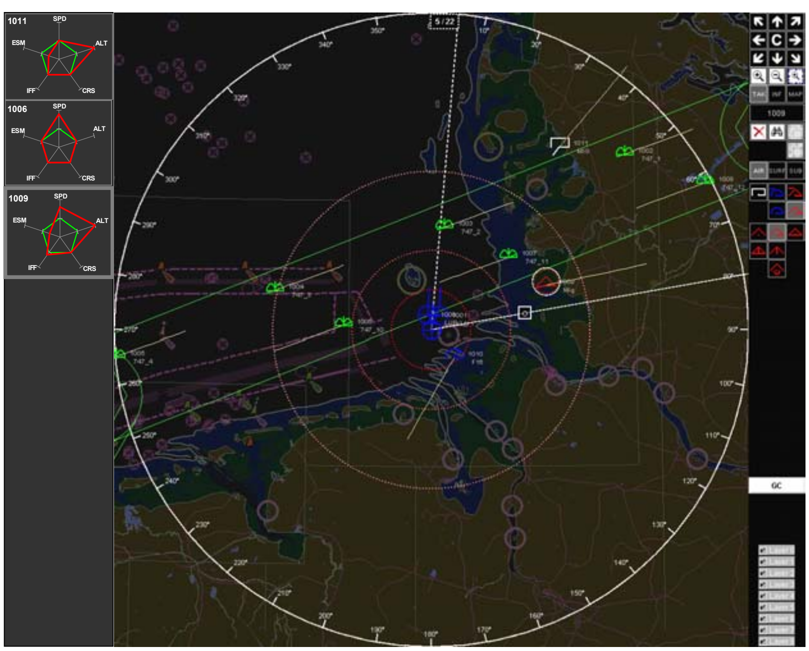

Figure C-2: Prototype of the Tactical Situation Display (TSD).

C.1.1.2 Management and Visualisation of Low Structured Data in Command and Control Information Systems

The work has been published in divers publications by Kaster, J. et al. (j.kaster@fgan.de)

The German military forces are more and more involved in international areas of conflicts for peace keeping, area control, evacuation operations as well as armed confrontation.

In the area of "Intelligence, Surveillance, Reconnaissance" (IRS) various different message formats and document types have to be processed in an integrated manner.

The following aspects have been considered:

- Information requirements for intelligence (G2);
- Activity flows while evaluating message contents;
- Processing functions in situation assessment; and
- Visualisation requirements for military situation representation.

Not only the tactical-operational scenarios are of importance, instead the recording, communication and evaluation of findings about military and non-military organisations become more and more important and relevant, whose structures, activities and developments have to be recognized.

A prototype has been developed that efficiently supports the various steps in the area of IRS:

- Input:
  - External sources provide various data and documents.
  - These are collected in a filing department.
- Administration and Analysis:
  - Managing the data;
  - Processing the data; and
  - Evaluating the data (inquiries).
- Output:
  - Transferring results to the headquarters.

This prototype (DBEins) is an integrated solution for structured acquisition, process and evaluation of messages in the area of military intelligence. It has demonstrated its usefulness and efficiency in various applications (military exercises, SFOR/KFOR, ISAF).

Features:

- Document management in a database: independent data storage and access.
- Electronic keyword index for effective sort and filter functionality for later evaluation.
- Support of free inquiries/full text index.
- Access control.
- Representation of the intelligence/message contents on a geographical background.

Visualisation examples as applications of DBEins by means of software agents (Figure C-3 – C-6):

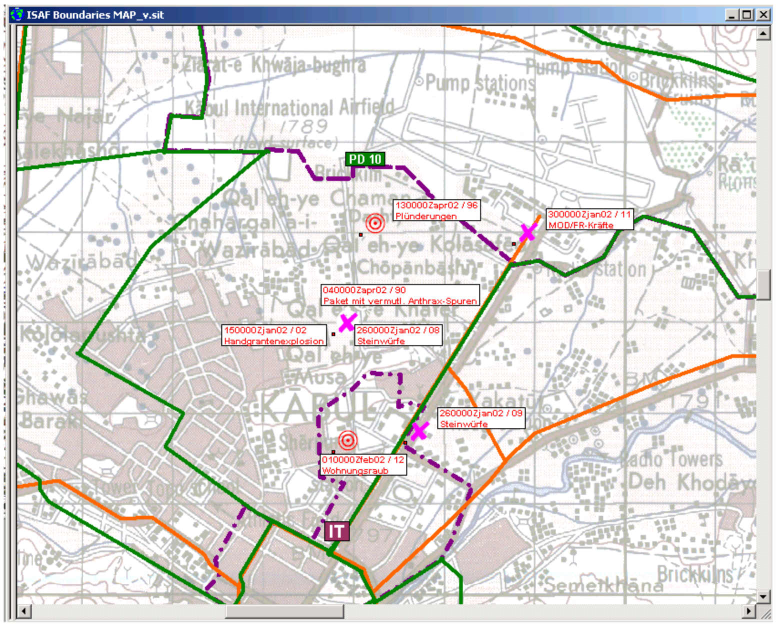

Figure C-3: Hotspots Kabul (Special Geo-Referenced Events, e.g. Crimes (Criminal Encroachments), Assaults).

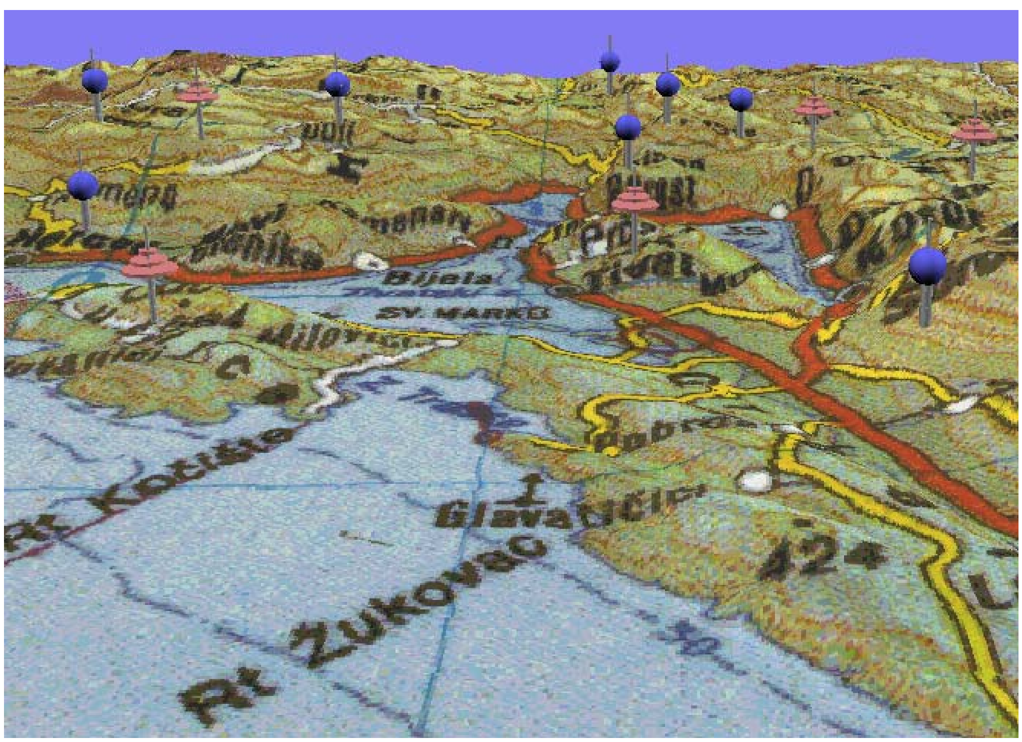

Figure C-4: Representation in 3D: (Interactive Exploration, Simulated Overflight).

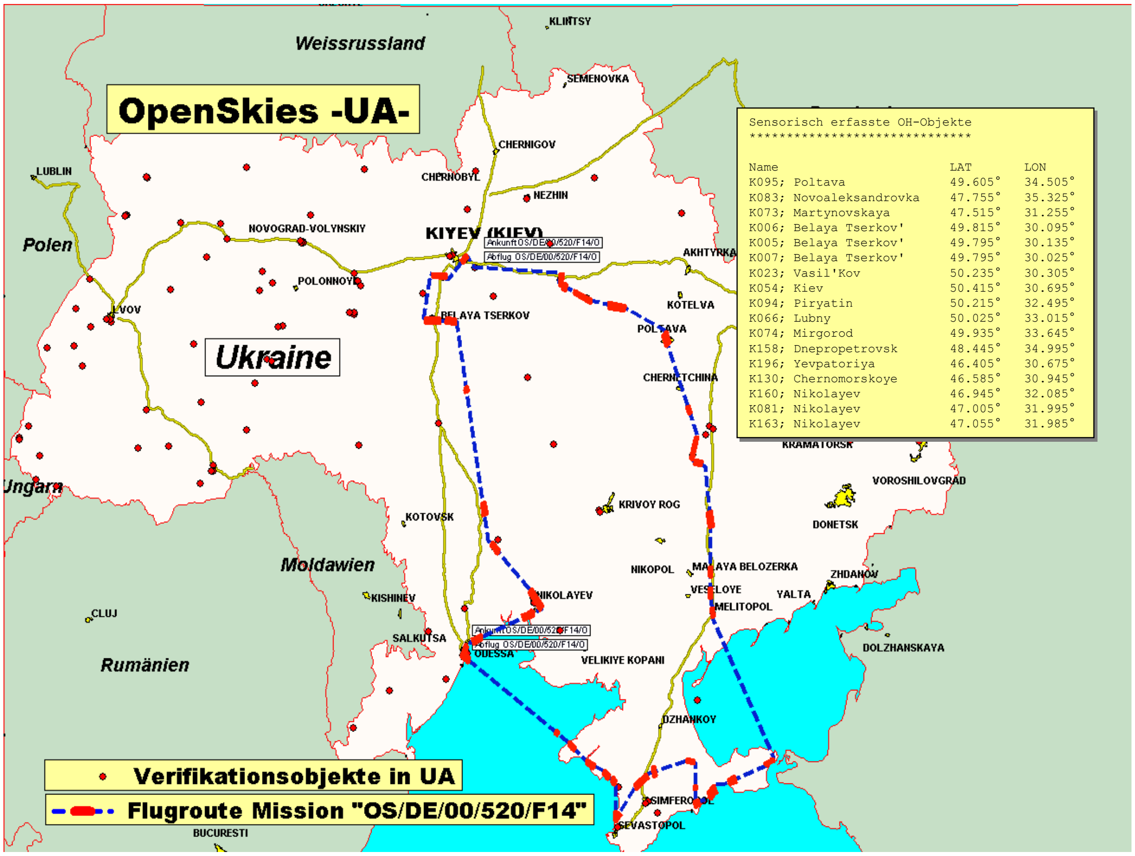

Figure C-5: **OpenSkies: By Superposing Geo-Referenced Data (e.g. Verification Objects) from Different Sources with an Air-Route Geographical Interrelations Become Transparent => Recognition, if and which Objects Have Been Elucidated During a Mission – A Report-Agent Provides a Textual Listing of the Results.**

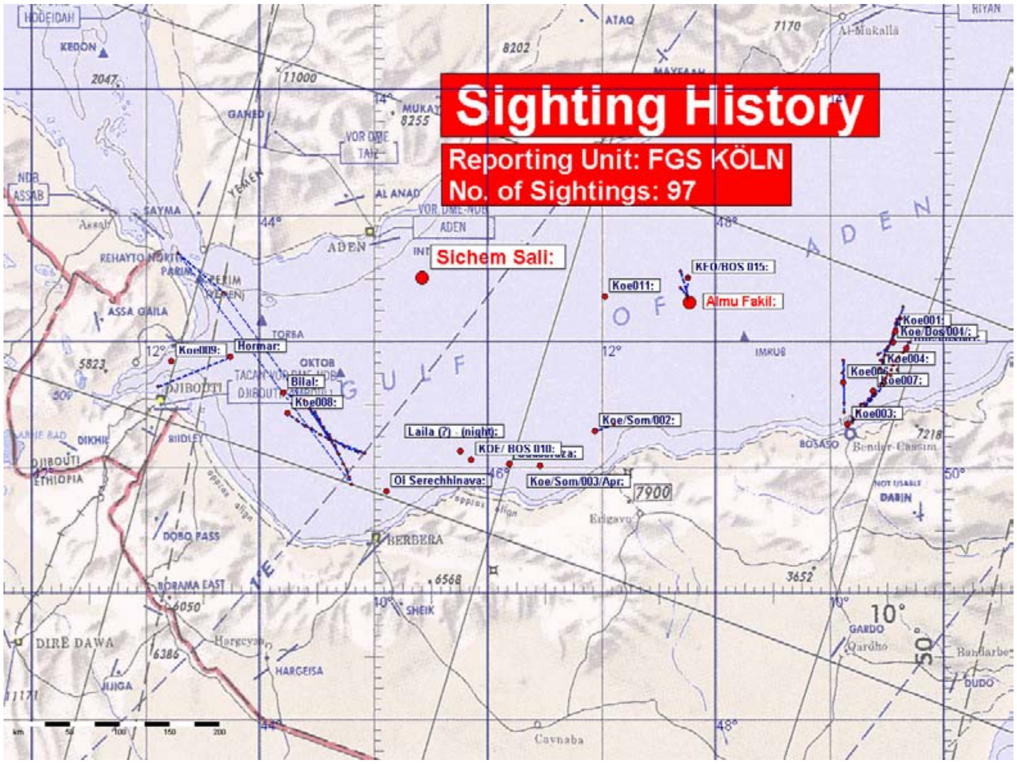

Figure C-6: Sighting History: Representation of Elucidated Ship Positions (Location: Horn of Africa, Operation "Enduring Freedom") – Not only Position, but also Track Information.

C.1.1.3 Evaluation and Design of AIS Information on ECDIS [2]

Empirical investigations were carried out in a research project for the German Federal Ministry of Transport, Building, and Housing to evaluate the presentation of AIS target information on ECDIS (Figure C-7). The investigations were performed at the Integrated Simulation Centre of Singapore (ISC) of the Maritime and Port Authority, at the Center for Marine Simulation (CMS) of the Memorial University of Newfoundland in Canada, and at the Maritime Simulation Centre (MSCW) in Germany. The features, colour and filling/size of AIS symbols, as well as the influence of the ECDIS display category on the detection of AIS targets were the main issues of the investigations. Results show that blue (S-52 colour token RESBL) is the most suitable colour of the tested colours for the presentation of AIS targets under all ambient light conditions on the tested IHO S-52 colour tables.

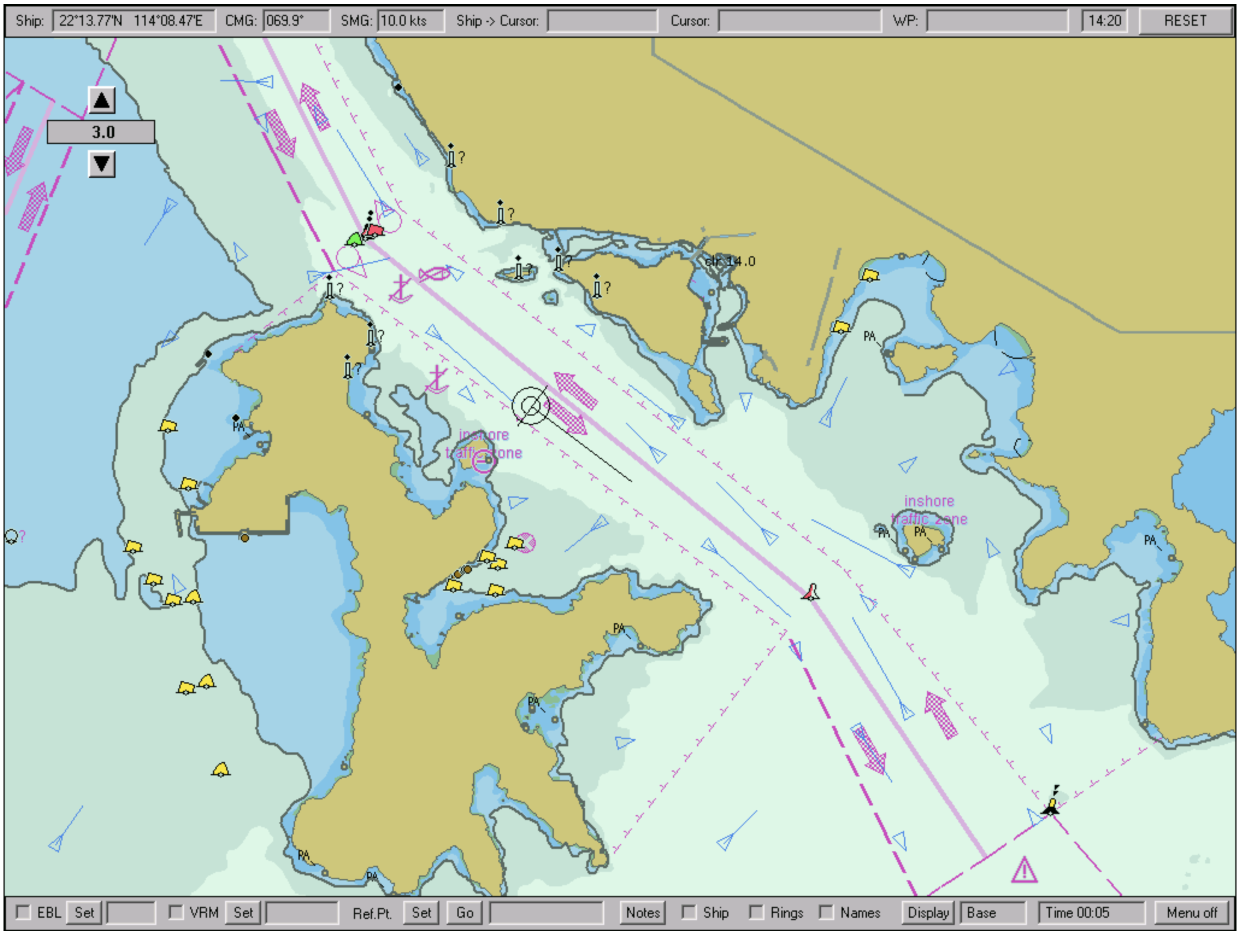

Figure C-7: Traffic Scenario on ECDIS (Electronic Chart Display and Information Systems).

C.1.1.4 Representation Concept for Conditioning the Combat Systems of the German Navy [3]

Parameterization of processes covered by modern software systems confers flexibility to systems operating in complex environments. Disadvantages of hard-coded systems like cost- and time-consuming software modifications no longer apply. Admittedly, complex multivariate systems or domains make high demands on the user due to the transfer of control functionality and responsibility and the task enhanced. Otherwise, the user may be supported domain specifically in several ways. Exemplary, knowledge- or rule-based planning components allow a situation dependent system configuration before a mission or specific parameters or combinations can be blocked as the situation requires.

There are many application domains that cannot be fully automated: for example, situations cannot be described or predicted, optimal adjustments cannot be determined either numerically or by simulation, the dynamic system is subject to permanent changes. In extreme cases the user must deal with all degrees of freedom, meaning an interaction of automated and manual adjustments is necessary. In these cases the user interface is of great importance. Not only the actual parameter values have to be clearly visualized but also the status of planning and support components in order to present and predict short-term future behavioral tendencies (system states).

The application domain examined is the combat direction system of the frigate class 124 of the German Navy. This is a semi-automatic system with a supporting rule-based control component like the hypothetical systems described above. This system comprises different functional processes with overall dependencies and provides degrees of freedom for the user, i.e. direct manipulation as well as automatic control are considered. A representation concept as well as an ergonomically designed user interface have been developed for the real mission as well as for integration in a simulation system.

The combat direction system of the frigate class 124 comprises the dependent intertwined processes sensor control, sensor data fusion, identification/classification, threat analysis, combat analysis, combat planning, combat execution and kill assessment (Figure C-8). The single processes are conditioned by means of operational parameters. These are parameters that can be manually adjusted by the operator or automatically by the use of doctrines. Doctrines are special decision rules that interlink actions and events and that are ideally defined as planning tools in the forefront of a mission.

Current research involves the formulation of a representation concept, the ergonomic representation of setting parameters and the design and development of an interactive user interface for directly adjusting the operational parameter. First, display formats of current parameter settings, orientations guides und access methods have to be developed that consider the diversity of parameter values and their logical groupings. Further, representation procedures have to be examined or generated for clarifying the coactions of various parameters as well as mission related dependencies. The user interface implies graphical representations of current system's adjustments and status of the rule base as well as edit options for parameter settings and designation of critical effects.

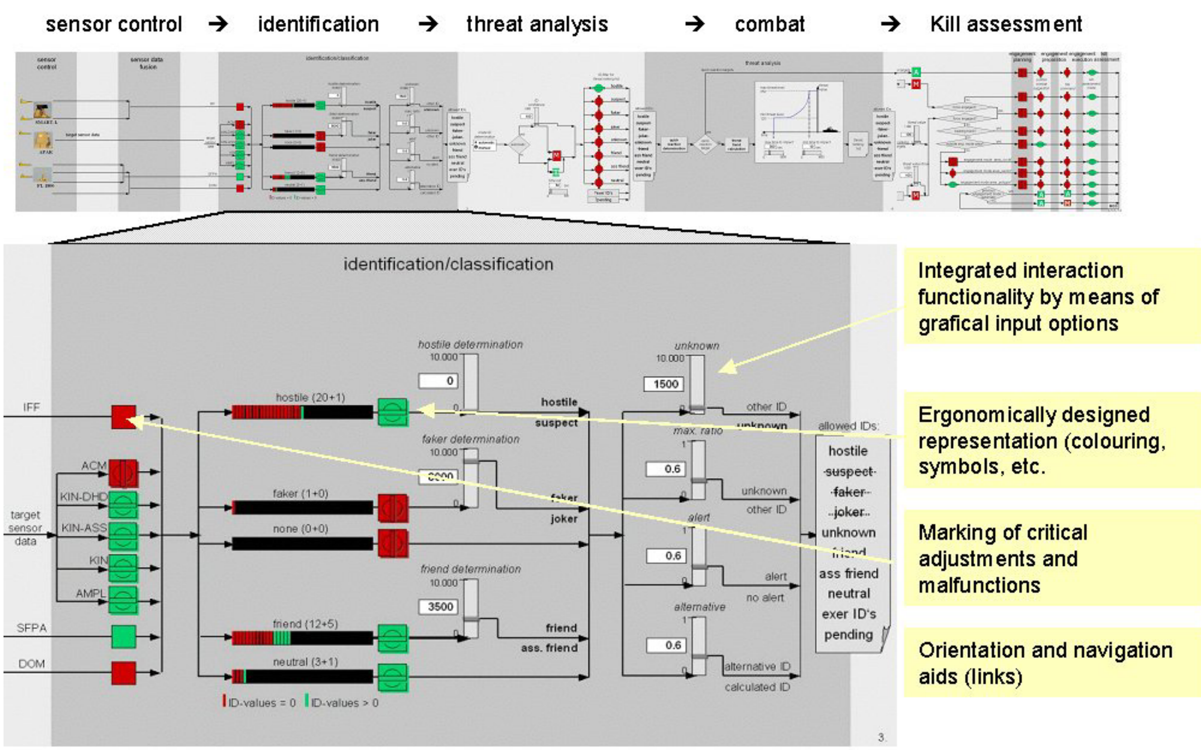

Figure C-8: User Interface for Manual Conditioning of Naval Combat Systems.

C.1.1.5 Visualising Relations (between People, Things, Cars, etc.)

In the area of Intelligence, Surveillance and Reconnaissance the identification, analysis, visualisation and evaluation of relations (areal, temporal and logical) are of great importance. For identification of the relations individual messages have to be analysed and relevant information have be filtered, as well as various messages have to be analysed and implicated.

In the view of actual counter terrorism aspects the logical/textual relations are of main importance, i.e. relations between people, between people and organisations, as well as relations between people and events, vehicles, etc.

The "DBEins" provides various ways to identify relations. Areal relations, e.g. can be represented/visualised by means of the GIS xIRIS (see above). Basis for the identification can be partial data structures, key wording of messages, linkages between documents or structuring of messages in partial databases.

Emphasis has been put on a widely automatic identification as well as the ergonomically designed representation of the relations.

First efforts have been made in identifying relations between people. The assumption has been made that people who appear in an individual message are related. The links between people in the People Database and the documents of the Messages Database have been automated. The evolving link list is the basis for visualising the relations. The modular design allows the integration of components based on computer linguistic in order to detect relations whose match is not obvious or assured ("Tom Smith" and "Thomas Smith" might be identical people, but "Tom Smith" and "Tom Smith" might not be).

Since the structure of relations is mainly network-like elements for representation like nodes and edges seem suitable. Like depicted in the picture nodes symbolize people, the edges symbolize the messages, which contain both people. The more edges between two people, the more messages deal about them, the surer is the relation.

The user selects interesting (or suspect) persons in the People-Database and obtains the graphical representation of the relations (Figure C-9).

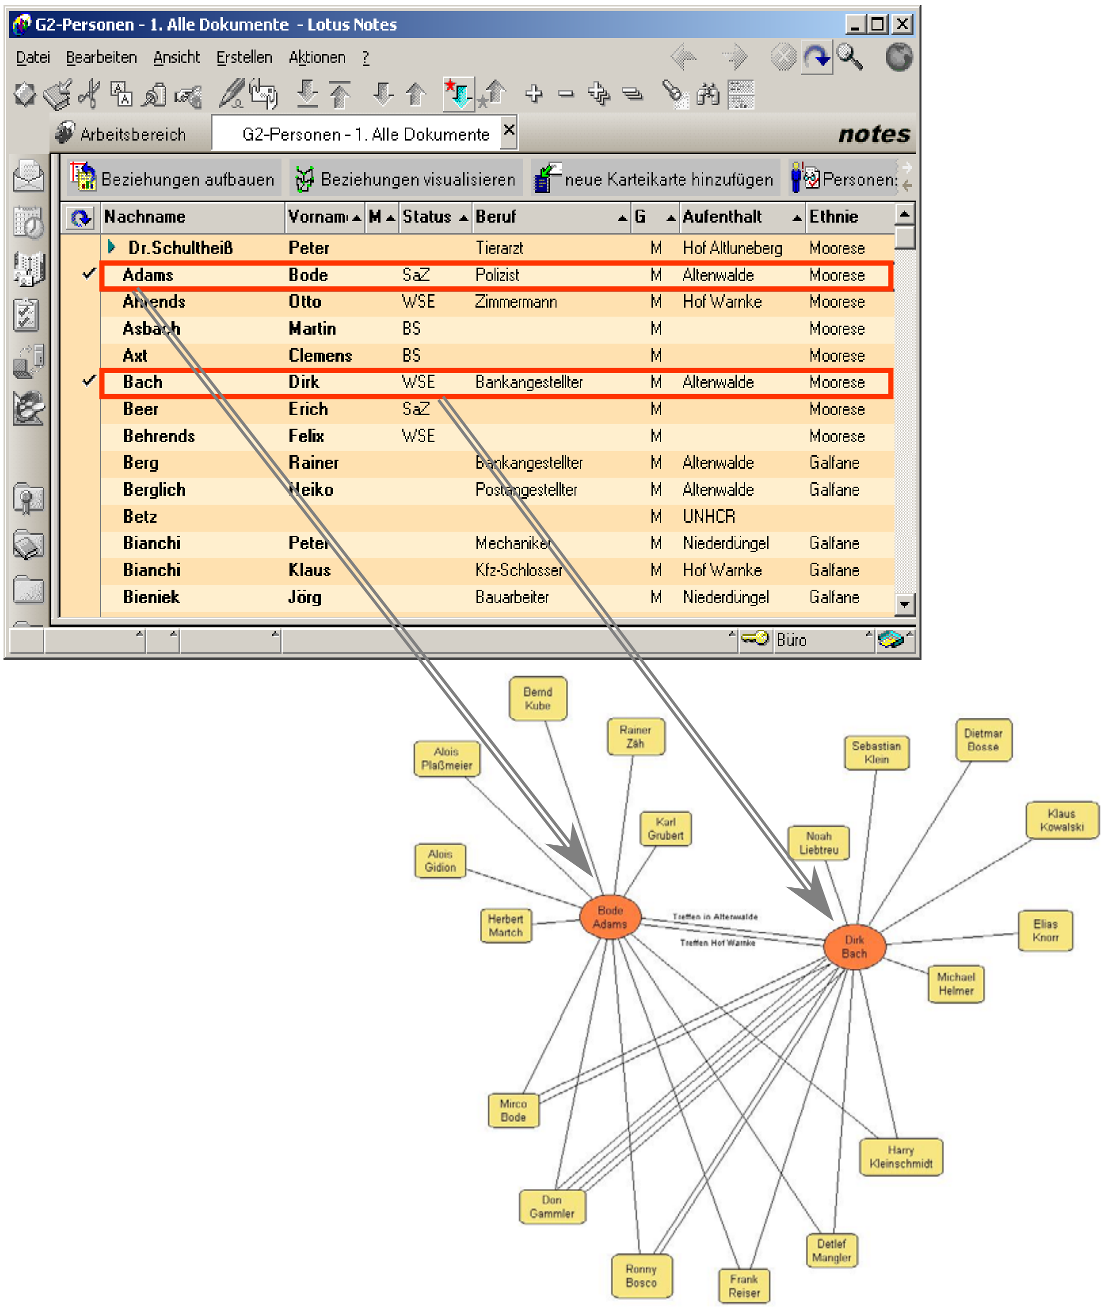

Figure C-9: People-DB and Graph of Direct Personal Relations.

C.1.1.6 Visualisation of Massive Amount of Data in Military Intelligence and Reconnaissance in Virtual Environments [4]

Due to the technology of modern information systems an increasing amount of data is available in military intelligence and reconnaissance. Today's software for supporting the user is not fully capable of coping with the new situation. Therefore the intelligence operator has to be supported by new, innovative methods for visualising the massive amounts of data. In this connection, the application of methods and technology of Virtual Environments is considered to bring along advantages for processing massive intelligence data. In collaboration with the project partners MEDAV, GmbH and Thales Systemtechnik GmbH a novel approach for visualising the complex information space of intelligence data and its various correlations within has been specified and implemented into a prototype. In a first step available approaches were summarized and analysed with respect to their practical usability. In a second step a prototypic interface between the data processing tool and the VE-framework of OpenActiveWrl of FKIE/EFS has been implemented. By this, technologies of VEs can be applied to an advanced exploration of massive data in military intelligence. (Figure C-10).

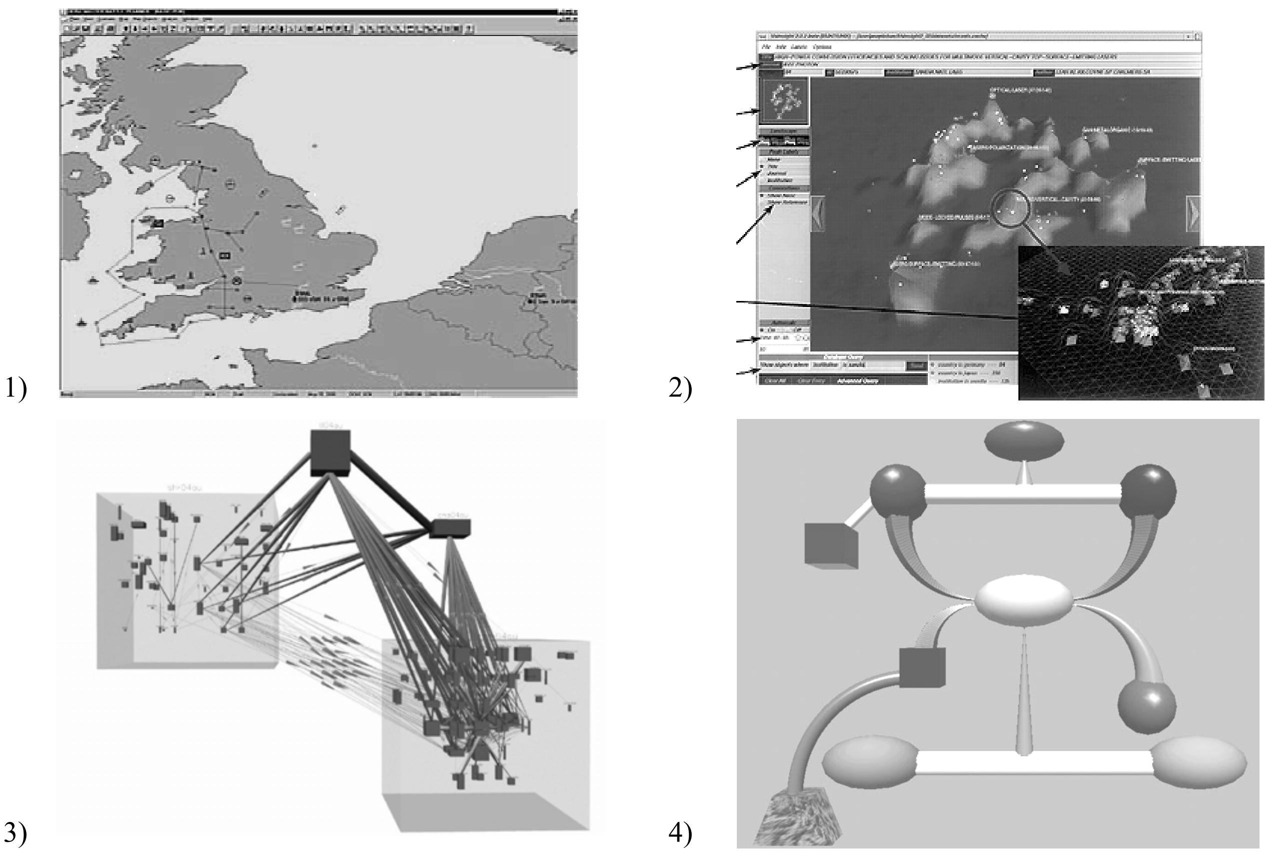

Figure C-10: Examples for Rudiments for Visualisation of Massive Datasets:

1) Visualisation of geographic and situation data (IST-013/RTG-002, 2001): Representation of military reconnaissance data assigned to specific locations. Communications data can be represented/visualised by connections/links between the reconnoitred locations.

2) Topography of a massive data set in documents (Sandia, 2002): Virtual topographies created by correlations within messages, not allowing geographical localisation, but clarifying semantically interrelations. Columns represent information clusters, valleys represent information sinks.

3) 3D-representation of relations between documents (NVision, 2002): Representing interrelations by means of a net. Nodes correspond to persons, terms, etc. and edges are defined by correlations.

4) Visualisation of data contents and their relations (VisLab, 2002): Like 3), but adding colour and shape to clarify similarities between nodes.

Representation of complex, massive datasets dependent on

- Mode and extent of the data to be represented;
- Mode and extent of the data to be visualised by the user/operator;
- Kind of output devices; and
- Target group of the operators.

Visualisation methods differentiated by

- Method of transformation of the abstract data into the **visual representation**;
- Method for reinforcement of effectiveness of the visual information (**reinforcement**);
- Adaptability of the representation to individual requirements (**interaction**); and
- Practicability of the visualisation solution (**operation**).

Methods

- Geographic referenced presentation of communications;
- Visualisation of abstract data sets;
- Graphics; and
- Completely abstract.

For the work conducted here visualisation will be twofold, i.e. geographic referenced representation for analogue communication nets and abstract representation for semantic correlations of message contents. Continuous transition between both representations is desirable for transferring contextual insights between both. The realistic, 3D and interactive presentation of VU (Virtual environments) provides high potential and innovative possibilities [5,6].

C.1.2 Fraunhofer Information and Communication Technology Group

The Fraunhofer Information- and Communication Technology Group (Fraunhofer IuK-Gruppe) represents the biggest, coordinated capacity in the field of applied research in informatics within Europe. The Fraunhofer ICT-Group hosts 14 institutes carrying out research in the area of information- and communication technology.

Fraunhofer IGD Rostock has four departments (Multimedia Communications, Innovative Interaction and Visualisation, Mobile Multimedia Technologies, Entertainment Technologies) and operates a Lab for Mobile Multimedia Applications, a Lab for Innovative Interaction Techniques, and a Lab for Multimedia and Telecommunication.

Information about the Institute may be found on the Web at:

http://www.igd-r.fraunhofer.de/IGD/index_html_en

http://www.igd-r.fraunhofer.de/IGD/Abteilungen/AR3/Projekte_AR3/ImVis/index_html_en

C.1.2.1 ImVis – Real Estate Visualisation with Cartographic Representation

The project "ImVis" is developed from the research concept "ActVis – Visualisation of activities", which deals with the representation of objects in the space time structure. To support the presentation and the sale of real estate in a better way, usually impressions from a helicopter flight will be shown, which includes the real property and vicinity. Due to high costs and weather dependency it became a mature technology to realize the flights virtual. However, so far these 3D-Animations only presents the real estate without showing the surrounding property and the nearer infrastructure.

In the research concept "ImVis – Real Estate Visualisation with cartographic representation" a system was developed, which includes – for the presentation of an object with textual information – both the location through cartographic representation and the surrounding property. At the same time it can be displayed – among the object of interest – additional information like the next stage of public transport, shopping facility, school/kindergarten, medical practitioner, entertainment and sport facilities, etc. stating the corresponding distance. The research concept "ImVis" is flexible and supports not only the 2D presentation, but it can be switched between the different presentation alternatives, e.g. video. Thereby it is possible to include and present location connected simulations, e.g. virtual flights of the department "Multimedia Information & Presentation".

C.1.3 University Paderborn

http://wwwcs.upb.de/fachbereich/AG/agdomik/forschung.html

Research on Collaborative Visualization & Augmented Reality, Web3D & Real-Time Shading is occurring at University Paderborn [7, 8].

C.1.4 University Bremen, Center of Complex Systems and Visualisation

http://www.cevis.uni-bremen.de/

CeVis, the Center of Complex Systems and Visualisation, is a research center at the Department of Mathematics and Computer Science of the University of Bremen, Germany. CeVis performs basic research in the areas analysis and visualisation of economic time series, cellular automata, chaos and fractals, wavelets, and the application of these concepts to various fields of science.

Furthermore, CeVis is involved in teacher enhancement activities in the USA and in Germany.

The origin of CeVis was the Institute for Dynamical Systems, which was founded in 1978. Its research group on Complex Dynamics formed CeVis in 1992. Starting as a basic research group working in the area of mathematics CeVis has developed to a research center with a focus on applications of the theories of chaos and fractals and wavelets to fields of research and development.

C.1.5 University Stuttgart

http://www.iste.uni-stuttgart.de/ps/Lehre/reengineering/Visualisierung.2.pdf

C.1.6 University Tübingen

http://www.gris.uni-tuebingen.de/areas/scivis/

C.1.7 University Bonn

http://web.informatik.uni-bonn.de/I/research/BibRelEx/Visualisation.html

C.1.7.1 Exploring Bibliographic Databases

For the exploration of bibliographic databases a visualisation system has been looked for that enables a 3D-representation of relations and comfortable navigation options. Knots and connections have to be selectable and provide detailed information about the source represented by this graphical element. The LEDA Platform of Combinatorial and Geometric Computing has been selected as the visualisation tool.

(http://www.igd.fhg.de/archive/1995_www95/papers/44/mukh/mukh.html)

C.1.8 University of the Saarland

http://www.st.cs.uni-sb.de/edu/einst/traces/

C.1.8.1 Collecting Trace-Data and Graphically Representing by Means of the Tool *aiSee* for Interactive Exploration

When working with any kind of complex relational data, visualisation provides for much better and faster understanding. Back in 1991, *aiSee* was developed to visualise the internal data structures typically found in compilers. Today it is widely used in many different areas:

- Genealogy (family trees, evolution diagrams, pedigrees);
- Business management (organization charts);
- Software development (call graphs, control flow graphs);
- Hardware design (circuit diagrams, finite state diagrams);
- Database management (entity relationship diagrams);
- Informatics (algorithm visualisation); and
- Webmastering (sitemaps, P2P networks).

aiSee reads a textual, easy-to-read and easy-to-learn graph specification and automatically calculates a customizable graph layout. This layout is then displayed, and can be interactively explored, printed and exported to various graphic formats.

aiSee has been optimized to handle huge graphs automatically generated by applications. It is available for Windows, Linux, Solaris, and Mac OS X.

C.1.9 University Konstanz

http://www.mpi-fg-koeln.mpg.de/%7Elk/netvis/trade/WorldTrade.html

A gallery of social structures; network analysis and network visualisation.

Figure C-11 is an example of social networks, depicting the volume and structure of world trade among OECD nations for the year 1992.

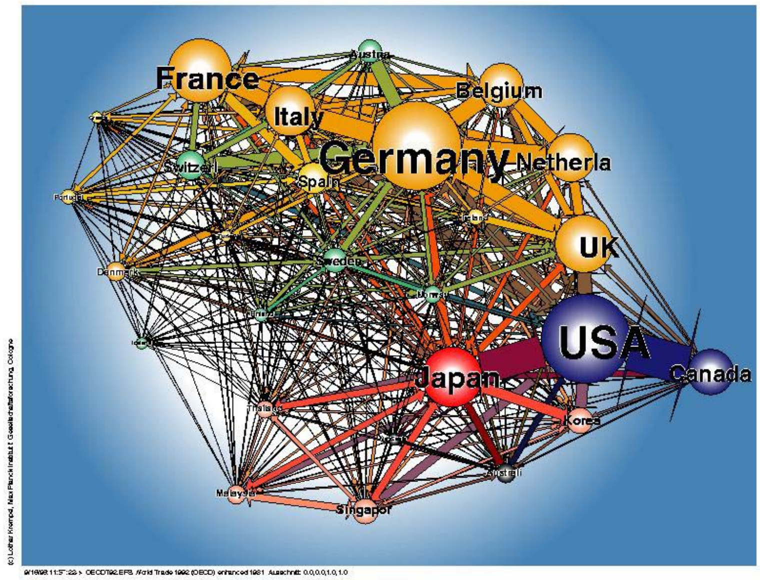

Figure C-11: Structure of World Trade in 1992.

C.1.10 Universität Heidelberg

http://www.vtk.org/

C.1.10.1 The Visualisation ToolKit (VTK)

The Visualisation ToolKit (VTK) is an open source, freely available software system for 3D computer graphics, image processing, and visualisation used by thousands of researchers and developers around the world.

C.1.11 Universität Bremen/CeVis

http://www.cevis.uni-bremen.de

CeVis, the Center of Complex Systems and Visualisation, is a research center at the Department of Mathematics and Computer Science of the University of Bremen, Germany. CeVis performs basic research in the areas analysis and visualisation of economic time series, cellular automata, chaos and fractals, wavelets, and the application of these concepts to various fields of science. Furthermore, CeVis is involved in teacher enhancement activities in the USA and in Germany.

C.2 REFERENCES

[1] Grandt, M., Distelmaier, H. and Pfendler, C., A Knowledge-based Human-Machine Interface for Future Naval Combat Direction Systems. *Proceedings of IST-043 Workshop: Visualisation and the Common Operational Picture*, Toronto, CA, 14-17 September, 2004.

[2] Motz, F., Widdel, H., Oei, P., MacKinnon, S. and Alexander, L. (2004). Investigations for Ergonomic Presentation of AIS Symbols on ECDIS. *International Hydrographic Review,* 5(2) (pp. 26-36). The Netherlands: GITC.

[3] Kaster, A. and Witt, O. (2004). Visualisation Concept for Operational Parameter Settings of Dependent Processes on Naval Vessels, Proceedings of IST-043 Workshop: Visualisation and the Common Operational Picture, Toronto, CA, 14-17 September, 2004.

[4] Visualisierung von Meldungsdatenaufkommen in Virtuellen Umgebungen, FKI/EFS/01/2004 Technischer Bericht, July 2004.

[5] Alexander, T., Goldberg, S., Delleman, N., Magee, L., Smith, E., McIntyre, H. and Worm, A. (2003). Compendium on Activities in Augmented, Mixed and Virtual Environments within NATO countries. Neuilly-sur-Seine, December, NATO RTA.

[6] Alexander, T. and Conradi, J. (2003). State-of-the-Art Report on Augmented, Mixed and Virtual Environments. FKIE/EFS/03/2004, Technical Report, November, FGAN. FKIE, Germany.

[7] Goetz, F. and Domik, G. (2004). openVisaar – Enriching cooperative visualization by combining open technologies, 1. International Conference on Cooperative Design, Visualization & Engineering, Palma de Mallorca, Spain, September.

[8] Goetz, F. and Domik, G. (2003). A Framework for Video-based and Hardware-Accelerated Remote 3D-Visualization, Eurographics 2003 Short Presentations, Granada, Spain, September, pp. 263-269.

Annex D – VISUALISATION PROJECTS IN DENMARK

Lisbeth M. Rasmussen
Danish Defence Research Establishment

D.1 INTRODUCTION

Visualisation-projects in Denmark can be categorised as either civilian or military. Civilian projects are either commercial or research and development (R&D). They will be mentioned in the next chapter. Military VR-projects are either simulators or R&D. The military projects will be mentioned in the following chapter. The final chapter gives a summary of the VR-projects in Denmark.

A common event for both the civilian and military society is the founding in May 1999 of DK-VRS (Danish Virtual Reality Society). DK-VRS still exists, but its activities have stopped (for further information, see http://www.dk-vrs.dk/).

D.2 CIVILIAN PROJECTS

Commercially visualisation is used by television stations, the entertainment industries and advertising agencies. Visualisation in research and development is of growing interest. Three universities have established VR laboratories or centres. These will be mentioned in the next section. Examples of research and development visualisation-projects will be mentioned in the following section.

D.2.1 Civilian Research Laboratories

D.2.1.1 Centre for Pervasive Computing

Pervasive computing is the next generation computing environments with information and communication technology everywhere, for everyone, at all times. Information and communication technology will be an integrated part of our environments: from toys, milk cartons and desktops to cars, factories and whole city areas – with integrated processors, sensors, and actuators connected via high-speed networks and combined with new visualisation devices ranging from projections directly into the eye to large panorama displays. The Centre for Pervasive Computing (CfPC) contributes to the development of

- New concepts, technologies, products and services;
- Innovative interaction between universities and companies; and
- A strong future basis for educating IT specialists.

Most of the work in the centre is organised as Research Projects involving both companies and universities. The 3D visualisation and interaction projects and the sound projects will be mentioned later.

Further information on CfPC can be found at http://www.pervasive.dk/index.html.

D.2.1.2 Centre for Advanced Visualisation and Interaction

CAVI (Centre for Advanced Visualisation and Interaction) is one of the research areas in the CfPC (see previous section).

3D visualisation is becoming increasingly widespread in as diverse areas as industrial design, architecture, city planning, medicine, moving images as well as the arts. The ability to interact in new ways with 3D models offers new possibilities for the professionals in these areas.

In a unique combination, the following 3D visualisation technologies are available in CAVI:

- 3D Panorama Cinema, Curved screen, Active stereo glasses, Tracking

  The Panorama cinema is a cylinder shaped screen placed in a room that seats approximately 15-20 persons. The size and shape of the screen mean that the visual angle of the spectators is almost covered by the screen. Models are displayed in active stereo, i.e. a picture for the right and the left eye is displayed alternatively in a very high frequency. Without shutter glasses the spectator experiences the image as blurred – but with shutter glasses that alternate between closing off for the right and the left eye in the same frequency, the spectator experiences an illusion of 3-dimensional depth. The Panorama is particularly suited for displaying large-scale models within the domains of architecture and city planning.

Figure D-1: 3D Panorama Cinema at CAVI in Use.

- TAN Holobench

  The Holobench is a combination of two 180x110 cm sized projection screens placed at right angles to each other in an L shape. Also here it is possible to show models in active stereo, i.e. the pictures on the screen draw a picture for the right and the left eye alternatively in a very high frequency. Without shutter glasses the spectator experiences the image as blurred – but with shutter glasses that alternate between closing off for the right and the left eye in the same frequency, the spectator experiences an illusion of 3 dimensional depths. A tracking system allows the spectator's gaze to be responded to through movement of the model and this heightens the impression of an object in front of the spectator's eyes.

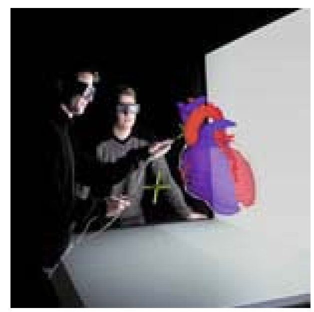

Figure D-2: Holobench at CAVI.

- Virtual studio

  The virtual studio or virtual stage set makes it possible to use digital 3D models as sets instead of scenography made from wood, steel, cardboard or other materials. Live recordings in a blue studio with real objects can be mixed with computerised models. It is, for example, possible for a studio host to walk into a completely blue room, sit down on a primitive wooden box and start reporting. What one sees on the television screen could then be that the host enters from a doorway, passes behind a row of plants to sit down at a speaker's desk to start reporting.

Figure D-3: Virtual Studio at CAVI.

- Performance Lab

  The performance lab is an open space and a facility for several types of experiments including experiments using a reactive performance space, which is a theatrical environment that enables physical actions to effect and manipulate electronic media. These spaces allow performers to improvise with media through a variety of means.

  Electronic media consists of any media that can be controlled from a computer. These are generally divided into four categories: visuals, light, sound and mechanical systems. Physical actions within the space consist of anything that can be sensed and interpreted by a computer. This consists of things like video based sensing, tracking systems, sound sampling, pitch detection or analogue sensors (heat, touch, bend, acceleration, etc.). CAVI has a motion capture equipment as part of the laboratory. Motion capture is a technology that makes it possible to register the movements of a person and use the data to animate a digitally created figure. Cartoonists use motion capture equipment to animate their characters. The technology is also used in many other areas for example in the study of movement in dance, sport and medical research.

Figure D-4: Performance Lab at CAVI.

- Onyx2 Infinite Reality2 Rack, 6X250 MHZ MIPS R10000, 1.5 Gb RAM, 2 graphics pipelines (each 64 Mb texture memory).
- Polhemus FASTRAK (3 sensors).
- Stereographic glasses.
- 6 Octanes.
- 20-30 O2.

Several CAVI-projects will be mentioned in a Section D.2.2.1.

For further information on CAVI see http://www.pervasive.dk/resAreas/CAVI/CAVI_summary.htm.

D.2.1.3 VR Media Lab

VR Media Lab (previously VR Centre North) is located at the University of Aalborg. Its main feature is:

- Cave with 6 walls, Active stereo glasses

  The Cave is a room of 2.5 x 2.5 x 2.5 meters in which continuous images can be projected onto side walls, floor, and ceiling. This creates a complete spatial presentation of the scene/model being shown, which gives the viewer a total "immersion" into the spatial virtual environment when using the "active" stereo glasses. By using a so-called electromagnetic tracking system the viewer can move in (or around) the visualised object. This installation is preferably for one viewer only and is designed for research and design development.

Figure D-5: CAVE at VR Media Lab in Use.

- 3D Auditorium, Passive stereo glasses

  The 3D Auditorium accommodates up to 80 persons placed traditionally in from of a large screen measuring 8 x 2.85 meters. The 3D Auditorium shows computer graphics in both 2D and 3D. It is possible to use the so-called “passive” stereo glasses thus presenting the objects as spatial, i.e. objects are seen as if the are in front of the screen. A magnetic tracking system has also been installed making interaction with the graphics possible. Besides the computer graphic images from other sources can be projected (PCs, VHS video and DVD), when the auditorium is used for traditional presentations.

Figure D-6: Powerwall at VR Media Lab in Use.

- Panoramic screen, Active stereo glasses, Magnetic-tracking device

  The Panorama accommodates up to 28 persons placed in front of a large cylindrical screen with a diameter of 7.1 meters, 160 degrees and a height of 3.5 meters. Because of the shape of the screen, the viewers are given a convincing spatial presentation of the virtual world. As in the 3D auditorium the Panorama can present computer graphics in both 2D and 3D as well as images from PCs, VHS video and DVD. Stereo visualisation augments the spatial effect, but the number of viewers is reduced. In Panorama a magnetic tracking system has been installed, and the experience is even more realistic as the screen to an even higher extent "surrounds" the viewers.

Figure D-7: Panoramic Screen at VR Media Lab in Use.

These are run by:

- Onyx2 Infinite Reality2 Rack, 6 graphics pipelines

  The three installations are driven by on large supercomputer, an ONYX2 IR2 from SGI. The ONYX is installed in a specially cooled engine room together with several other larger servers. VR Media Labs ONYX2 IR2 has 16 parallel CPUs, 2 GB Ram, 288 GB HD in Raid3 system (transfers 100 MB/sec.) and 6 graphic pipes, etc.

- PC cluster

  A PC cluster is being established.

For further information on VR Media Lab see http://www.vrmedialab.dk/pr/index_e.html.

D.2.1.4 Centre for 3D GeoInformation

Centre for 3D GeoInformation (3DGI) brings together research, public authorities, and business communities in a unique environment of developing 3D GeoInformation applications, all based on new Virtual Reality technology as well as on information regarding urban and rural areas. This is done by establishing a Virtual Geographic Infrastructure (VGI), enabling a wide range of geographically related information to be spread via new, netbased means of communication. One of the new aspects in the project is the user interface, based on intensive use of Virtual Reality (VR) and 3D. By creating a virtual 3D model of reality and then use it as an index for many other types of information, it becomes possible to use the general human ability to familiarize with the surroundings and navigate through space.

The goal is to establish a pioneering project, which will be the central force for the very latest within VR and GIS technologies.

The purpose of Centre for 3DGI is gathering knowledge and competence during the process of creating 3D models of cities and landscapes for organising and presenting GeoInformation applications.

This will be done by:

- Collecting competence and knowledge within the field by arranging seminars/conferences, establishing international research networks and by employing researchers within this particular field;
- Collaborating with companies, who already possess the most recent competence within VR and 3D urban and rural models or are interested in acquiring this;
- Establishing a VR user interface for looking for position-fixed information in the northern part of Jutland;
- Creating a geographical model of North Jutland, which can form the basis of digital visualisation and the marketing of the resources of the region;
- By developing a basis of knowledge and documentation for the use of a geographical communication concept covering the northern part of Jutland, adapted to the expected increased band width in digital transmission medias (Fixed and Mobile Nets) and as a framework for developing virtual environments; and
- Forming the basis for future research and for building up regional knowledge within field-gis (field registration with mobile units). augmented reality (a mixture of 3D models and reality), three-dimensional user interface and the use of broad band for mobile knowledge services.

Further information on Centre for 3DGI can be found on http://www.3dgi.dk/en/3dgi.html.

D.2.1.5 VR•C

VR•C is a VR centre at the Technical University (DTU) in Lyngby north of Copenhagen. It is collaboration between UNI•C (a national IT-centre under the ministry of education) and The Technical University (DTU).

The objective of VR•C is to further utilization, research and education in the field of virtual reality (VR) in Denmark and Scandinavia.

Primarily virtual reality is used in connection with building activities, architecture, design and research.

VR•C's main facilities are:

- Holowall (TAN Powerwall (6.5 m x 2.5 m)), 2 x 3 projectors (TANORAMA/Electrohome) for passive stereo, Ascension Flock of Birds Tracking system with 2 trackers + 1 6-degrees-of-freedom mouse, 8 active speakers, 3 subwoofers, 1 Microsoft SideWinder Forcefeedback Pro joystick, 1 control for audio panel and SGI terminal.

  In the Holowall as many as 50 people at a time can experience the virtual world. By means of stereo projection and polarized spectacles the 3D objects appear spatially in front of the screen. Used for building trade, product development and research-oriented projects.

Figure D-8: Powerwall.

- Holodeck, 3 n-vision DataVisor HiRes Head Mounted Display, See-through (augmented reality), 2xXSGA CRT, Headphones and microphone, Ascension MotionStar tracking system with 8 trackers, 3 x 24" Monitors.

  In the Holodeck the user wears a VR helmet and sensors on his/her hands to attain integration with the virtual world.

Figure D-9: HoloDeck.

- SGI Onyx2 InfiniteReality2 Graphic Supercomputer

  16 195 MHz MIPS R10.000 CPU's

  8 GB RAM

  356 GB Hard disc

  3 graphic pipes each with 2 Rastermanagers and up to 24 displays

- Lake Huron 20 sound computer

  8 Quad DSP processors

  32 I/O ports

- 4 SGi O2 Workstations

- 1 SGi 320 Personal Workstation

Keywords for VR•C are industry, education and research. Areas of interest are architecture and landscape modelling, scientific visualisation, visual simulation and education, and collaborative VR. For further information on VR•C see http://www.uni-c.dk/generelt/english/research/vr-c.html.

D.2.1.6 FORCE Technology

FORCE Technology merged in 2002 with DMI (Danish Maritime Institute).

FORCE Technology offers simulator facilities and tools for all levels of maritime education from computer-based training to full-mission simulation.

Figure D-10: View from the Simulator.

FORCE Technology possesses a number of full-mission training simulators – one with 360° graphics on 18 m diameter screens and all with a full range of bridge equipment. All can be operated individually or interactively with full passage communication between the ships.

For further information on FORCE Technology see http://www.force.dk/gb/default.htm.

D.2.2 Civilian Visualisation Research and Development Projects

D.2.2.1 Projects at Aarhus University

D.2.2.1.1 Virtual Urban Planning

Virtual urban planning (VUP) is a research project involving Aarhus University, Cadpeople and COWI. The project's goal is to develop and investigate the use of virtual reality technology, as a basis for improvement of decisions to be taken in a city region that is continually undergoing development. VUP has to function as a combined tool for the benefit of politicians, public administrations, building constructors and architects, private businesses and the town's citizens. A 3D-model is created on many levels of detail-addition, where future alterations to the physical relation in the town space can be inserted and tested. In this way, the model becomes dynamic. If, for example, there are several construction suggestions for the same project, they can be "turned on and off" in the model and the individual changes can be evaluated in a more extensive context, since the surroundings are created in 3D. After this, not only the building process, but also its relation to the surroundings can be evaluated.

Further information can be found at http://www.pervasive.dk/projects/virtUrbPlan/virtUrbPlan_summary.htm.

D.2.2.1.2 Visualisation of the Cadiovascular System

The cardiac morphology in patients with congenital heart disease is often very complex and variable from individual to individual. Consequently, accurate morphological information remains of outmost importance when planning the surgical intervention. Magnetic Resonance Imaging (MRI) is the imaging modality that currently provides the best soft tissue contrasts. New visualisation techniques based on three-dimensional MRI have been developed and is now being implemented and tested clinically at Aarhus University Hospital. The preliminary results show that cardiac morphology in congenital heart disease can be accurately reconstructed and represented by virtual models. As MR image quality continuously improves, image processing times are being reduced rapidly. This will undoubtedly make three-dimensional MRI with virtual reconstructions an important clinical investigative tool within the nearest future. Related research will implement elastic tissue properties, allow cutting the model and inserting patches, as first steps towards a surgical simulator/trainer.

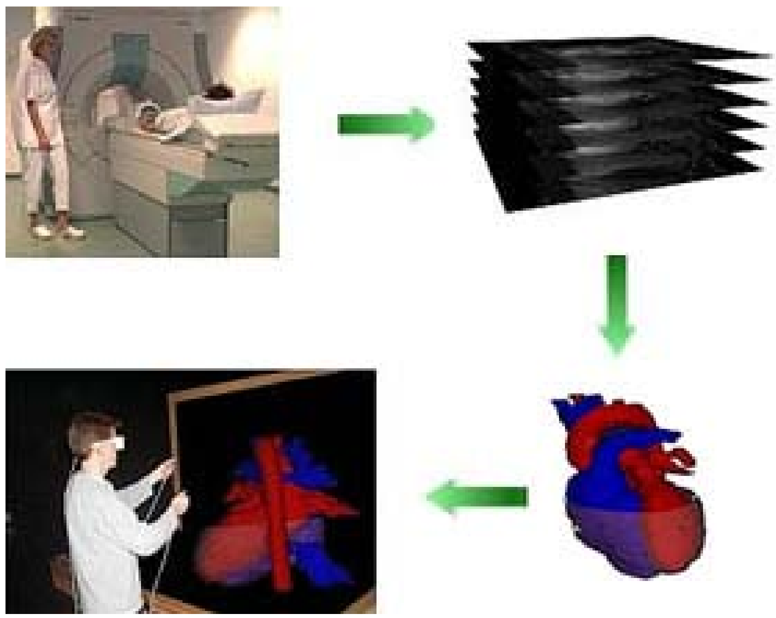

Figure D-11: The Process from Scanning to Visualisation on Virtual Reality Equipment.

Further information can be found at http://www.pervasive.dk/projects/cadiovas/cadiovas_summary.htm, Virtual Reality Heart (Systematic Software Engineering), and http://www.daimi.au.dk/~svorre/heart.html.

D.2.2.1.3 Digital Theatre – Hyperopticon

The use of digital media in theatre and other time based art forms is increasing as the cost of computers decreases and the development of software programmes has been made more available to theatre technicians. Still there are many fields still to be researched and explored. The members of The Digital Theatre from the University of Aarhus have worked in the following fields of digital theatre research: the dramaturgy of digital media, virtual puppet theatre, motion capture/animation, and reactive spaces. Through experiments and productions The Digital Theatre group has explored the production and reception of theatrical spaces in order to explore interactive possibilities in digital media.

Figure D-12: Digital Theatre.

HYPEROPTICON was created from a concept of developing a particular site i.e. in a library, where a small audience might have the opportunity to explore a specific theme, a play or an event through interaction in a mixture of digital media and real-life performance.

The scientific goal is to develop concepts and produce digital theatre experiments that can be used to further the knowledge of dramaturgical understanding and broaden the knowledge of perception of digital time based art.

One focus of the research has been to develop new forms of staging plays through an idea of exhibiting a story or a plot and to make this an interactive experience. So the investigation of space and digital technologies is also a crossover art experiment with digital media as means of production.

Further information can be found at http://www.pervasive.dk/projects/hyperO/hyperO_summary.htm.

D.2.2.1.4 Video Prototyping

Within experimental system development there is a long tradition that visions for future technologies are concretized by means of prototypes, mock-ups, scenarios and video. The virtual studio makes it possible to mix digital elements with persons and physical objects, thereby opening up new prospects of combining the above mentioned concretization possibilities in the form of video prototypes. An internationally known example of a visionary video is Appel's Knowledge Navigator from 1987 (see www.billzarchy.com/clips/clips_apple_nav.htm).

Through production and application of a number of video prototypes the purpose of the project is to investigate their application in the communication of visions for future technologies, and in design workshops, etc.

The results are video prototypes within the following areas:

- Handheld display;
- The future hospital ward;
- The virtual architect;

- The future drawing office;
- The future school; and
- Living and interactive buildings.

Figure D-13: Video Prototyping.

Further information can be found at http://www.pervasive.dk/projects/vidProt/vidProt_summary.htm.

D.2.2.1.5 3D Experiences

3D opens up completely new avenues and sets entirely new demands on all aspects of filming, whether dramaturgy, production or staging. Also in the case of exhibitions, communication, entertainment and education, 3D technology is increasingly important. Today 3D is already in use in adventure parks, such as Legoland parks and Universal Studios, but 3D technology confronts great challenges and possibilities not only in this area but also in showrooms, exhibitions and education.

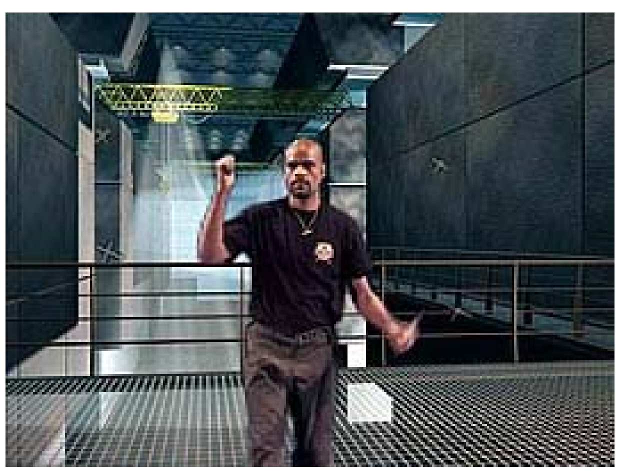

Figure D-14: Using VR for Staging.

The overriding aim is to investigate and develop the potential for 3D within film, interactive TV and other forms of sense-stimulating and educational areas. The research goals are:

- To reveal and develop present and future technological possibilities and challenges within the 3D narrative and experiences; and
- To investigate new possibilities and visions, which arise when classical film artistry meets virtual reality.

Further information can be found at http://www.pervasive.dk/projects/3Dexp/3Dexp_summary.htm.

D.2.2.1.6 3D Sound in 3D Space

3D visualisation in Panorama displays is a widespread means of communication within a number of areas including architectural-design and experience-oriented applications. At the same time, it has been confirmed that sound similarly provides a powerful instrument which so far has been exploited to a limited degree in 3D visualisation in panoramas and similar interactive 3D display facilities. It is the project's overriding aim to develop and investigate the use of 3D sound in three-dimensional graphical spatial models. The research goals are:

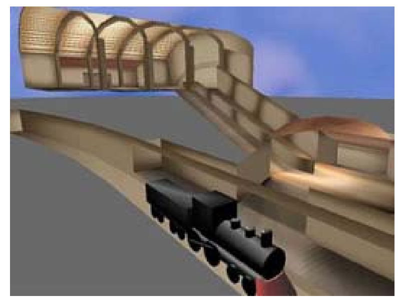

Figure D-15: 3D Sound in 3D Space.

- To develop the immersion experience in 3D presentations in Panorama using the audio dimension;
- To identify future potentialities for exploiting the audio dimension in conjunction with 3D visualisations within architecture, design and experience-oriented applications; and
- To improve simulations of the natural sound experience in an interactive environment.

Software will be developed to make it possible in connection with a visual 3D model to move around in it, where 8 sound sources are positioned in the model and the sound reception corresponds to the sound position and orientation with respect to the sound sources and where the sound experience reflects the 3D room's acoustic properties. In this way the sound can be "coloured" so that it reproduces the virtual environment's sound-reflective properties more faithfully, and thereby strengthen the immersion experience, which the panorama provides.

An algorithm will be developed, whereby the sound can find its way to the listener in the virtual world. That is, the sound's path to the listener takes the height of the walls, doors, corridors, etc., into account.

Further information can be found at http://www.pervasive.dk/projects/3Dsound/3Dsound_summary.htm.

D.2.2.1.7 Interfaces for 3D Applications

3D models are becoming increasingly widespread and used in a diversity of areas, for instance architecture, product design, medicine, and art. At the same time pervasive computing brings us into the domain where software and 3D models must be portable between systems with different display and interaction systems some of which are integral part of physical devises or other parts of the environment. Whereas studies and development of interaction technology and styles for 2D user interfaces have a longstanding tradition, navigation and interaction in a 3D environment is fairly unexplored.

The scientific goals are to advance the development of cross-platform independent interface technology for 3D virtual reality applications, and to develop a universal non-touch interface for navigation and interaction in a 3D environment.

Further information can be found at http://www.pervasive.dk/projects/interf3D/interf3D_summary.htm.

D.2.2.1.8 Sound as Media, between Signal and Music

Sound is used in our every day artifacts as signals e.g. when an alarm goes on, and as jingles to make us aware of e.g. the news on TV. Sound is also used to accompany computer games and other multimedia products. And finally, sound is used in its own right as music. Hence, sound can be used functionally and/or emotionally. Sound can be generated by very cheap devices so that sound signal can be included in all our artifacts, and sound can be generated on large scale computers and played through very sophisticated layouts of speakers in physical and virtual spaces. Sound design will hence play a major role in the design of physical and virtual spaces massively equipped with digital technology.

Figure D-16: Sound Experiment.

Further information can be found at http://www.pervasive.dk/resAreas/Sound/Sound_summary.htm.

D.2.2.1.9 Distributed Multimedia Technologies and Applications (DMM)

The DMM project focuses on research issues central to the development and use of distributed multimedia. Distributed multimedia technologies possess much potential for improvement of work-practices and collaboration in network organisations, e.g. through intranets; in society, e.g. electronic commerce and mass communication; and in distributed education, e.g. virtual seminars and distance learning. To explore this potential, the project will undertake three main kinds of experimental research activities:

- Participatory analysis and design in the three listed application domains;
- Technology research and development within the areas of collaborative hypermedia, virtual reality, and real-time multi-modal communication; and
- Demonstrator development, pilot tests, and evaluation in an industrial network organisation as well as an educational setting.

With respect to network organisations and electronic commerce, the empirical and the demonstrator activities will take place in modern Danish network organisation applying multimedia technologies both in products, commerce, and business processes. With respect to distributed education, the corresponding activities will take place in the context of establishing joint multimedia courses at Masters and Ph.D. levels between the Technical University of Denmark and Aarhus University. The results of the project will include theories about distributed multimedia, prototypes of new, distributed multimedia technologies, and documented user experience.

Further information can be found on http://www.intermedia.au.dk/projekter/dmm/index.html.

D.2.2.1.10 Finished Projects

- **3D image processing for cranium- and brain-surgical planning and simulation** at CAVI and PET Centre at Aarhus University Hospital.

- **Digital, 3D atlas of the receptor systems of the human brain** at CAVI and PET Centre at Aarhus University Hospital.

- **Product Development**

  CAVI has for 2 years cooperated with Centre for Product Development at the Technical University of Copenhagen on visualisation. The following are examples where 3D visualisation has been used:

Table D-1: Finished Product Development Projects at Aarhus University

Arla Foods	Packaging design
Kampsax/Lundbeck	Landscape visualisation
Arkitektfirmaet Schütze A/S	Architecture
CF Møller/CADpeople	City planning
Danmarks Radio	Virtual set activities / Ren Kagemand (Danish television show)
Jydsk Dykkerfirma (Diving firm in Jutland)	Beaching Museum St. George

- **The Family Factory**

  The family factory is a theatre show that combines ordinary theatre with traditional animation, puppet handling and live 3D computer animations. The virtual creations are not just programmed ahead with a limited set of actions. Actors control them with motion capture while they perform their one role. The project is a corporation between CAVI, The Danish Film School, and Schule für Schauspiel-kunst "Ernst Bush" Berlin.

  Further information can be found on http://www.multimedia.au.dk/JCal/ff/fabrikinfo.html**.**

- **The Digital Theatre**

  An Experimentarium in 1999 examined the artistic possibilities involved in the encounter between real and virtual performers, between virtual beings in real spaces and real beings in virtual spaces. These encounters were established by various means of digital technology transforming movements in time and space into 3D-animations via digital data. The technology used for this purpose was motion capture technique, animation programs and projection techniques. Further information can be found on http://www.daimi.au.dk/~sdela/dte/index.html.

- **Whizbang**

  From a set of measurements of the seabed, a data file is created which is visualised by means of volume rendering. It is possible to cut in the large amount of data and manipulate the model in different ways so that the data is easy to grasp. A Master's thesis project by Niels Husted and Kaare Bøgh.

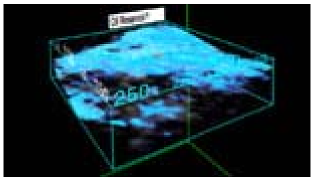

Figure D-17: Whizbang.

- **Aarhus new art museum**

  A model of Aarhus's new art museum for the panorama cinema. The model was made in co-operation with the art department Schmidt, Hammer and Lassen for the Aarhus Art Museum and financed by the merchant department of the Aarhus City Council.

Figure D-18: Aarhus New Art Museum.

- **The CAVI building**

  A digital model of CAVI's building which was used in connection with planning and designing the building.

Figure D-19: The CAVI Building.

- **Katrinebjerg Phase 2**

  A model of the Katrinebjerg Phase 2 building was made to show visitors to CAVI how the next phase of the IT park extension will look.

Figure D-20: Katrinebjerg Phase 2.

- **Katrinebjerg Phase 1**

  A 3D model showing an overall vision of the future Katrinebjerg IT City.

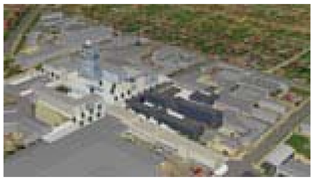

Figure D-21: Katrinebjerg Phase 1.

- **Architectural competition**

  Visualisation of a project proposal for the TDC domicile made for the architecture firms of Jørn Schütze.

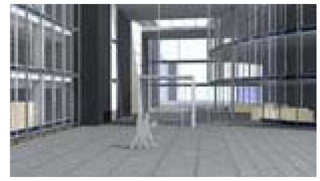

Figure D-22: Architectural Competition.

D.2.2.2 Projects at Aalborg University

D.2.2.2.1 The Project of Sonderborg

VR Media Lab has in collaboration with Cadpeople and COWI made a 3D model of some parts of the town of Sonderborg. The project was financed by the Danish National Research and Educational Buildings in connection to a new university building project.

Further information can be found on (page is under construction) http://www.vrmedialab.dk/pr/activities/spatialmodeling/sonderborg.html.

D.2.2.2.2 Aalborg University Campus Model

VR Media Lab is building a 3D model of the campus at Aalborg University. The model is going to be used as tool for the next 10 to 20 years when the university is going to be enlarged.

Further information can be found on (page is under construction) http://www.vrmedialab.dk/pr/activities/spatialmodeling/campus.html.

D.2.2.2.3 CAE and CFD into Virtual Reality

Computer-Aided Engineering (CAE) and Computational Fluid Dynamics (CFD) into Virtual Reality (CCVR) is a post-doc planning research activity at VR Media Lab by Truc Huynh and Henrik R. Nagel. CAE and CFD are developed for structural engineering application, typically on the personal computer monitor. However, it is a common desire for structural engineers that a large screen is needed for large structure, where the whole construction can be analysed in global dimension or zoomed in into a local part as well. Further, the large screen and the virtual reality device will also make the idea visible during the research step and the final presentation. The CCVR research is to the effect the creation and analysis of a construction in virtual reality (VR), where the suspension bridge is an example.

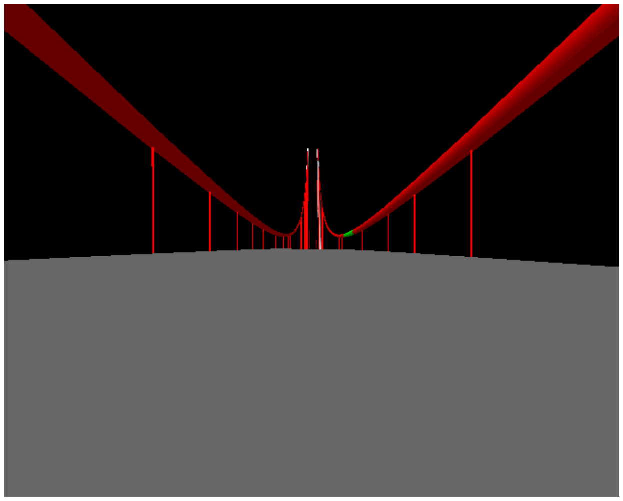

Figure D-23: CAE and VR in Bridge Construction.

Further information can be found on http://www.vrmedialab.dk/pr/activities/simulation/ccvr.html.

D.2.2.2.4 3D Airflow

VR Media Lab has in co-operation with the Department of Building Technology and Structural Engineering (also at Aalborg University) worked on visualising airflows in buildings.

Aalborg University, the Danish Institute of Agricultural Sciences, and the Royal Veterinary and Agricultural University have concluded a five year framework program together focussing on computer simulation of ventilation, airflow, and indoor climate of stables. The purpose of the project is to understand and control the airflow in stables. This will make it easier to improve the working conditions for the farmers, the animal welfare, and the economy of the animal production.

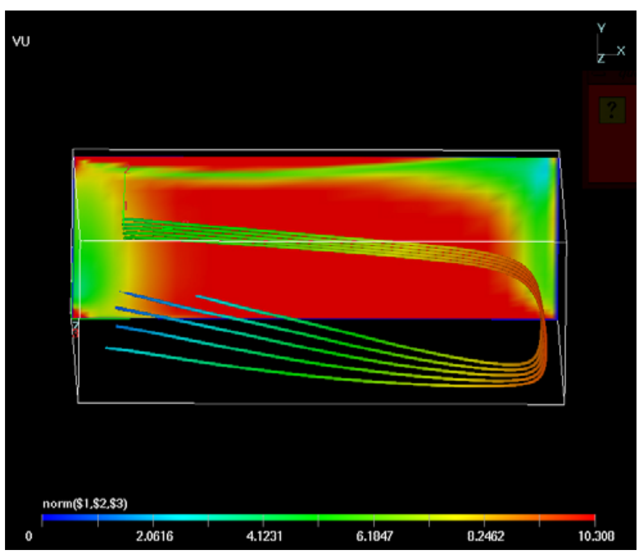

Figure D-24: 3D Airflow in a Stable.

The role of VR Media Lab's is to visualise the 3D computer simulations being run at the Royal Veterinary and Agricultural University and at the Department of Building Technology and Structural Engineering.

The purpose of the visualisations is to examine the details of the results of the computer simulations, which can be difficult to perceive on ordinary computer screens. The results have also been shown to consultants and companies in the ventilation business.

Further information can be found, in Danish, at: http://www.vrmedialab.dk/pr/aktiviteter/simulering/3dluft.html; an English version is expected to be available soon at: http://www.vrmedialab.dk/pr/activities/simulation/3dluft.html).

D.2.2.2.5 Data Mining

Data mining is a research project at VR Media Lab. The goal of data mining is to obtain useful knowledge from large collections of data. The discovered knowledge can be rules describing properties of the data, frequently occurring patterns, clusterings of the objects in the database, etc. Current technology makes it fairly easy to collect data, but data analysis tends to be slow and expensive. In most standard database operations nearly everything the user sees is something that they knew existed in the database already. Data mining extracts information from databases, which the users did not already know about.

Visual data mining is visual, explorative analysis of large amounts of data. The starting point is data without the any hypotheses about the data. The process of analysing the data involves interactive, usually undirected search for structures, trends, etc., and the result is visualisations of the data, which provides the users hypotheses about the data.

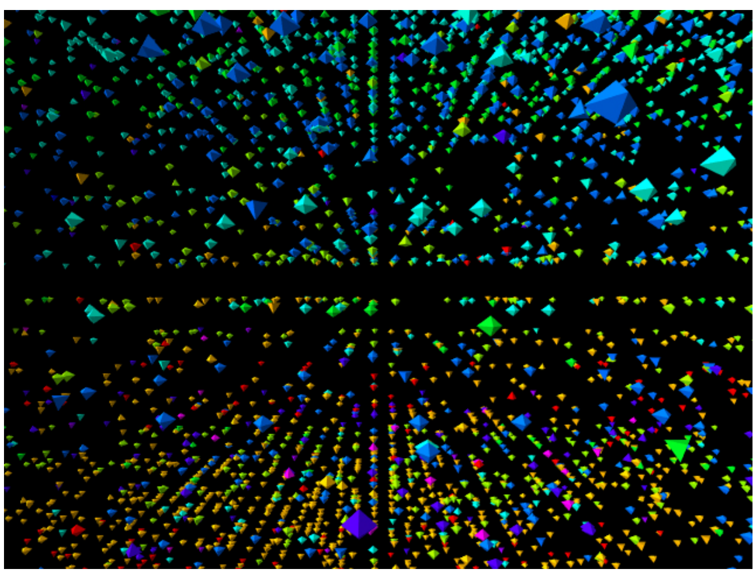

Figure D-25: Data Mining.

Further information can be found on http://www.vrmedialab.dk/pr/activities/datamining/overview_d.html.

D.2.2.2.6 3D Visual Data Mining

3D Visual Data Mining (3DVDM) is a research project at VR Media Lab.

Technology to store and process large amounts of data has during the last decades improved dramatically. This has led many companies to store ever increasing amounts of customer information in large databases. The hope has been that it would be possible to discover unknown relationships in the data, and thereby obtain a knowledge which could give commercial advantages. However, finding hidden relationships in large amounts of data is not easy. Purely numerical methods have been supplemented with visual methods. This has led to the emergence of "Visual Data Mining". Visual Data Mining has traditionally employed 2D techniques, such as geometric, icon-based, pixel oriented, hierarchical, and graph-based methods. With the 3D Virtual Reality (VR) facilities available at VR Media Lab it is now possible to explore how the ability to interpret 3D objects can be used in Visual Data Mining. 3DVDM aims at exploiting all the possibilities that immersive 3D VR technology may provide to Visual Data Mining, and then adapt the data processing methods to this. The project has collaborators from four areas of science: Database Systems, Statistical Analysis, Perception Psychology, and Visualisation. Researchers in the database area are responsible for handling and delivering large amounts of data from databases. Statisticians are responsible for finding interesting problems, creating statistical models for data analysis, and guiding the entire project in experiments of 3DVDM system. Perception psychology researchers are responsible for finding ways of good interaction and human perception of VR environment. Visualisation researchers are dealing with the visualisation of data using VR Media Lab facilities and implementation of a novel VR environment for data mining.

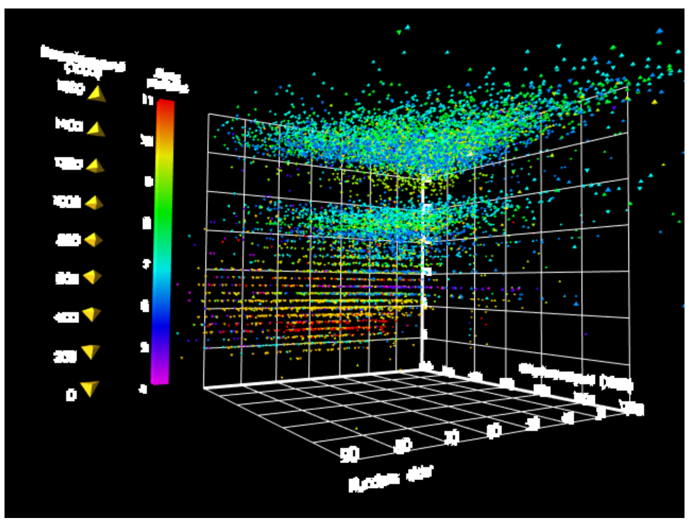

Figure D-26: Visual Data Mining.

Further information can be found at http://www.cs.auc.dk/3DVDM.

D.2.2.2.7 Learning

VR Media Lab is developing a pilot project (in cooperation with Vestas and Zenaria) in order to clarify the use of VR-technology in various learning situations. This particular pilot project involves teaching of Vestas electricians thus enabling them to perform a more accurate repair of the windmill on-site. This, however, requires a more thorough understanding of the theories behind basic electromagnetics. The problems arise from the fact that foreign electricians do not have the same educational background as Danish electricians. Therefore the outcome of the project is an education, where the windmill electricians attend a course at VR Media Labs facilities followed by an introduction to a CD-ROM which contains VR-sequences from the course.

Further information can be found on http://www.vrmedialab.dk/pr/activities/learning/overview_l.html.

D.2.2.2.8 Vestas

One of the learning projects at VR Media Lab is called Vestas. The project addresses problems at two levels: at an abstract and method developing level as well as at a concrete application level. At the first level the project presents the following problems:

- To include the newest forms of interactive multimedia and VR technology in order to develop and test new methods of learning in a virtual environment.
- To estimate to what extent different interactive technologies can be used in learning situations regarding with education and further training.
- On the basis of a specific education program to evaluate the use of educational methods and possibilities.

At the second level the following problems are presented:

- To develop a VR product for Danish as well as foreign wind turbine electricians in order to extend their qualifications.

The wind turbine industry has a considerable amount of export trade. The mounting and the following maintenance are done by foreign electricians who may not be fully qualified. Besides there may be cultural differences causing an incorrect translation of the Danish manual.

Figure D-27: Visualising in the Vesta.

An incorrect mounting or maintenance may lead to a breakdown of the wind turbine causing unnecessary financial costs. This project aims at solving the above problems by using VR technology.

By using visual means the project will give the necessary information thus being able to cope with any language or education related barriers.

Further information can be found on http://www.vrmedialab.dk/pr/activities/learning/vestas.html.

D.2.2.2.9 Finished Projects

- **A Music House in Aalborg**

  The city of Aalborg needs a new music house ("Musikkens Hus"). The new music house will be situated by the harbour right next to the old power plant.

  VR Media Lab made a visualisation of the proposals, which were sent from several architectural firms from all over the world. These visualisations were presented to the committee of judges. Their decision was made based on the visualisation of each proposal.

  The winning proposal was presented at "Studenterhuset" (student house) in Aalborg in February 2003. Furthermore a report in collaboration with the Danish National Research and Educational Buildings has been issued. This report describes the entire project using the VR technology in an architectural competition.

  Further information can be found on http://www.vrmedialab.dk/pr/activities/spatialmodeling/musikkenshus.html.

- **A plan for new housing development**

  The project of the town of Ans, Kjellerup Borough in mid Jutland is an ordinary plan for new housing development. These housing plans are made every year all over Denmark. The characteristics of this particular project are that the new housing area is at the outskirts of the town. The new housing area is situated in attractive natural scenery overlooking a meadow.

Figure D-28: The Town of Ans.

COWI, a Consulting Engineers and Planners firm in Aalborg, asked VR Media Lab if they were interested in visualising this new housing development, as the borough of Kjellerup was very sceptical about the extent of the plan and the damage done to the area. The first step was to model the project area. COWI prepared the basic data from altitude information, ortho photos, and from a road project cut into the model.

VR Media Lab employed four students from the Architecture and Design study program, who prepared a first draft to the borough during the summer of 2000.

The next step was an alteration in the regional plan. The model was modified according to the alterations agreed upon, and the project was presented to all members of the Kjellerup Town Council in May 2001. Based on this presentation it was decided that the project was to be implemented. The area has been site developed and is now ready for its new residents.

For further information see http://www.vrmedialab.dk/pr/activities/spatialmodeling/ans.html.

- **Highway at Holbæk**

The project aimed at testing a large road project based on the facilities and equipment of VR Media Lab.

Model conversion, model reduction, and navigation were at the top of list of things to be tested. A very simple model was presented. Only the most necessary parts of the project were modelled; many spatial elements in the landscape – e.g. buildings, plants, road elements – were not included in the model. Right now there are no plans for continuing this project, even though the Panorama is an obvious facility to present landscape constructions of this particular kind.

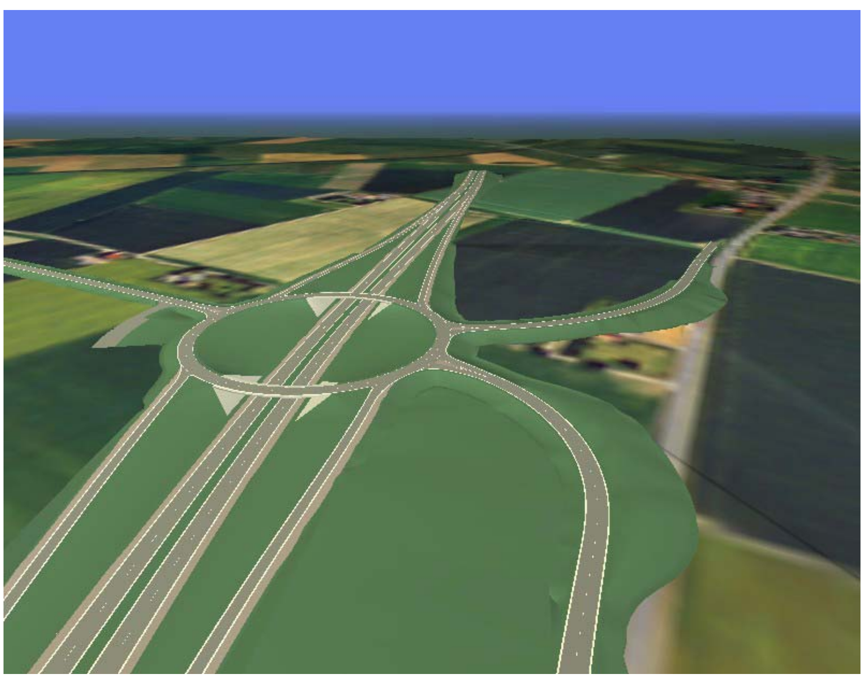

Figure D-29: Highway at Holbæk.

For further information see http://www.vrmedialab.dk/pr/activities/spatialmodeling/motorvej.html.

- **Model of NOVI 3, 4, and 5**

  The project aimed at modeling the NOVI buildings. This enabled a presentation of the models in the Panorama facility causing a discussion about the further building plans of NOVI.

  The project was divided into two parts: The first part deals with NOVI 3 and 4. This part was completed in the summer of 2000. Simultaneously models of NOVI 1 and NOVI 2 have been outlined.

  The model of NOVI5 was made at a very early stage compared to the construction itself, which was completed in the summer of 2001. VR Media Lab worked together with architect Peter Tybro from Vilhelm Lauritzen AS.

  The architect firm made traditional drawings of the building, and the modeling has been made from two-dimensional technical drawings of ground level and aspects.

  The model was being used to describe various light angles at different times of the year. These light angles were presented and evaluated in the Panorama facility. The model was also used to present choice of interior, colours and forms inside the building.

Figure D-30: NOVI.

For further information see http://www.vrmedialab.dk/pr/activities/spatialmodeling/novi.html.

- **Puppet** was a research project funded by EU-commission at the University of Aalborg. It involved developing virtual inhabited 3D rooms for educating pre-school children.

- **Staging of Virtual Inhabited 3D-spaces** was a research project funded by the Danish National Research Council involving the University of Aalborg, etc. It dealt with all kinds of aspects regarding the nature and usage of the signs system of interactive multimedia; which in general terms meant the semiotic of interactive multimedia systems. The purpose was to define a universal common language for 3D interactive multimedia systems.

- **Sound in Cyberspace** is a research project at the University of Aalborg.

D.2.2.3 Projects at Centre for 3DGI

The projects at Centre for 3DGI can be divided into the following Research/Development Tasks:

- Automated Extraction of 3D objects;
- Qualification of 3D GeoInformation;
- Queries for navigation in 3D Models;
- Representation of objects;
- Distributed Database System; and
- Viewer.

Further information on 3DGI can be found on http://www.3dgi.dk/en/3dgi.html.

D.2.2.3.1 Qualification of 3D GeoInformation

3D models of landscapes based on ortho photos and laser scannings is becoming accessible on the internet and in other 3D environments. The technique of the software is advancing rapidly, but not much guidance for building up models optimal for the users exists. The goal of the research group of 3DGI is therefore not only to develop 3D geographic software, but also to take user demands into consideration, both concerning graphical design of objects and data input, and the functional design like navigation and interaction. Developing interfaces, design and navigation at 3DGI is based on user tests and inspiration from classic and more well-established research areas like cartography and web usability/GUIs.

Further information can be found on http://www.3dgi.dk/en/research/qualification.html.

D.2.2.3.2 Queries for Navigation in 3D Models

The main research focus lays on queries related to the 3D Geographical database, i.e. a repository containing data related to the surface of the earth. The query result will be a pre-processed data set suitable for 3D visualisation. The aim of the queries is to enable an efficient navigation through huge amount of data stored in one or more interconnected databases. The amount of data transferred to a viewer should be minimized to potentially displayed data only. The underlying query algorithms will take into account the level of detail of the geographical features. It will also perform automatic data pre-processing, regarded as cartographic generalization in 3D. Only few working algorithms have been introduced, which deal solely with 2D cartographic maps. A similar approach combined with the level of details could result in generating a more meaningful 3D model, i.e. avoiding unnecessary details, while graphically emphasizing the features of interest.

Further information can be found on http://www.3dgi.dk/en/research/queries.html.

D.2.2.3.3 3DGI Viewer

The 3DGI viewer uses part of the ROAM algorithm to do level of detail on the terrain. The algorithm uses fewer triangles to represent the flat areas and more to represent the uneven areas like hills, cars, trees, and houses. The images shown in the viewer are dependent on the viewpoint dependent version of the viewer. Not only does the viewer use fewer triangles to represent the uneven areas, now it also uses more triangles for the structures close to the viewer and reduces the number of triangles far away, but is also uses viewpoint dependent continuous level of detail on the landscape. Landscape patches are geo referenced, and the viewer can handle more than one. A segmentation process of the laser scanning groups points that belong to the same plane in the same segment. In particular this is noticeably on the larger segments such as the ground and the rooftops. The first rooftops have been reconstructed. The algorithm is still unstable and makes assumptions about the houses. Click on picture to compare the segmentation images and view rooftop.

Figure D-31: 3DGI Viewer.

For further information contact info@3dgi.dk.

D.2.2.4 Research Projects at the Technical University in Lyngby

The following sections gives short presentation of research projects at VR•C and mentions a few finished projects from before the VR•C was established. For further information on the existing projects please contact vr-c@uni-c.dk.

D.2.2.4.1 Interactive 3D Visualisation of Projects

Interactive 3D visualisation is an area of growth where interactive virtual models are used to support decision-making in connection with large projects and design of new products. The virtual model is used to improve the level of information and communication. The virtual environments are based on CAD models and technical 3D drawings that are imported into the VR system thus enabling the user to navigate and work interactively with different 3D objects.

D.2.2.4.2 Scientific Visualisation

With scientific visualisation researchers have a tool for handling and visualising theoretical results of their research work. It is possible to obtain new insight by making abstract models and large complex amounts of data accessible and intelligible.

D.2.2.4.3 Finished Projects

- **Network-Based VR technology**

  Network-based VR technology is an innovative research area. With visualisation and broadband technology it is possible for groups across subject boundaries and geographical distances to collaborate on complex problems in a virtual 3D environment.

- **VR in Neuro Informatic** at the Technical University.

- Simulating a combine harvester at the Technical University.

D.2.2.5 Projects at COWI

COWI is involved with the Virtual Urban Planning project and the project of Sonderborg. COWI also has development projects by itself in the areas of GIS and VR. One of these is the Skyline project.

D.2.2.5.1 Skyline

COWI is distributing 3D visualisation software from Skyline Software Systems. With this software and the nationwide 3D model a virtual flight across Denmark is possible. TerraExplorer (the Skyline 3D-viewer) makes it possible to freely navigate in a 3D world. The photo-realistic, aerial image based terrains can be accessed on the Internet through video streaming – without appreciable loss of performance. In TerraExplorer Pro all the desired terrain overlay information, such as routes or text, labels and graphics, to promote specific features in the landscape can be imported. Figure D-32 gives a representative snapshot from TerraExplorer.

Figure D-32: Aarhus, as shown using TerraExplorer.

Skyline project: COWI has supplied the County of North Jutland with a Skyline 3D-model of the county, composed of a height model, draped ortho photos and supplied with building polygons from The National Survey and Cadastre Department. Everybody can go into the North Jutland County homepage, www.3d.nja.dk, and fly around in the county or you can key in an address and fly to the address in question. Figure D-33 shows the simulator console for this project.

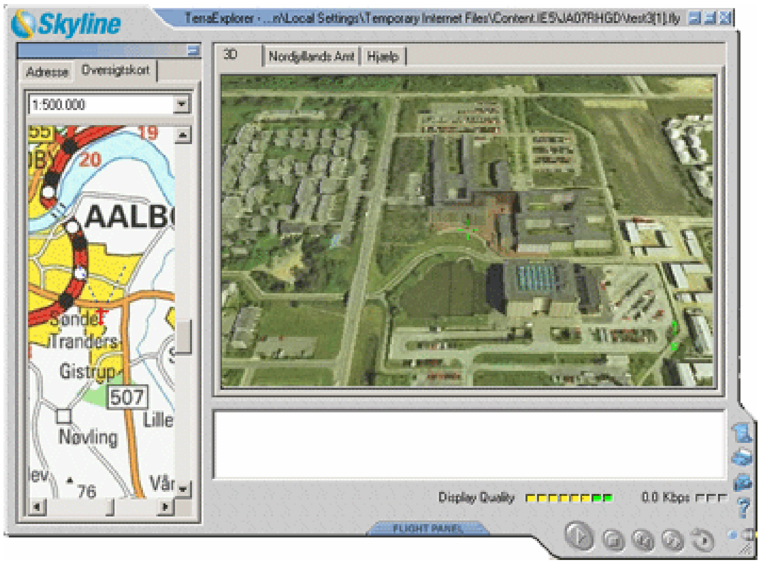

Figure D-33: Skyline Project for the County of North Jutland.

D.2.2.6 Finished Projects Elsewhere

- **Vision Quest** was a visual interactive form of education in natural sciences. It was a PhD project at UNI-C.
- **Digital 3D-architecturemodels on the Internet** was a pilot project at the Art Academy Architecture School.
- **Visualising the information in a library system** was a prototype at the School of Commerce.
- **Motion Capture** at 3D Connection.

D.3 MILITARY VISUALISATION-PROJECTS

Military Visualisation-projects are either ready-made bought simulators or research and development. Exceptions are Mission Debriefing System for the Air force, VIKING GIS, and RTP 6.14 Virtual Environment in Military Decision Making mentioned below. Military research and development takes place at the Danish Defence Research Establishment (http://www.ddre.dk/).

D.3.1 HCI-Lab

The HCI-Lab at the Danish Defence Research Establishment is used for testing VR hardware and software. In 1999 the laboratory had the following hardware:

- Teranetix Blackbird XL Xeon
  2 x 3.06GHz Xeon 533 FSB
  2 GB ECC DDR memory
  182 GB hard disk
  ATI FireGL X1 256 MB AGP Pro x8

- Assorted PCs
- Ascension's Flock of Birds
- V6 helmet
- I/O glasses
- Gloves
- Logitech Spacemouse

The laboratory has the following software:

- SuperScape
- Sense8
- Multigen Creator
- DI-Guy

D.3.2 Tactical Trainer for a Group Leader

The tactical trainer for a group leader is a research and development project at the Danish Defence Research Establishment. It started 1998, and is a tactical trainer for education and training a group leader. It combines a programmed simulation model with a geographical information system and a VR system and will have a speech interface. It will probably finish 2004. The first prototypes are 'Attack of mechanised infantry'. The conceptual model is pictured in Figure D-34.

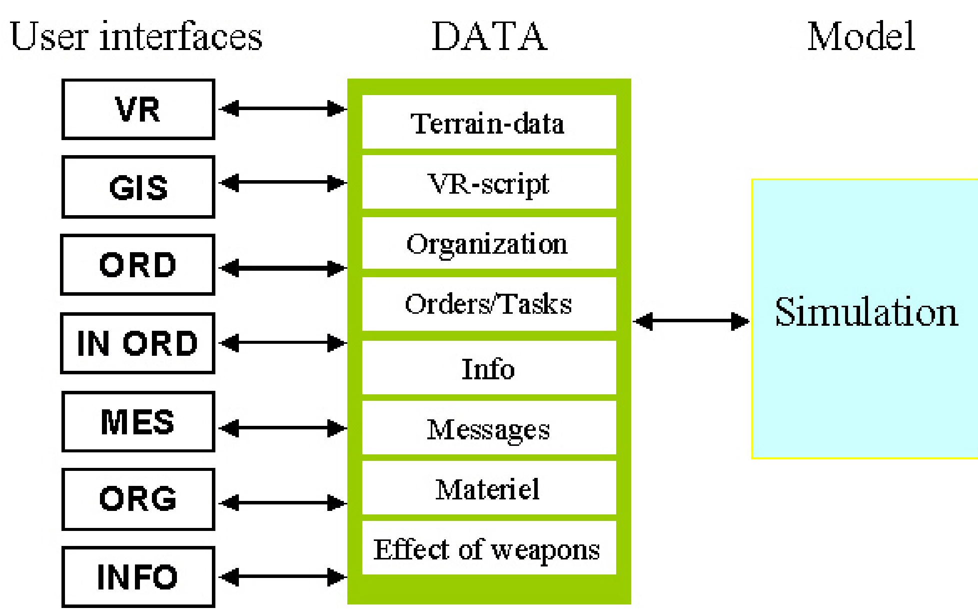

Figure D-34: Conceptual Model.

The ordering of the soldier is done by voice. The VR part was developed in Superscape.

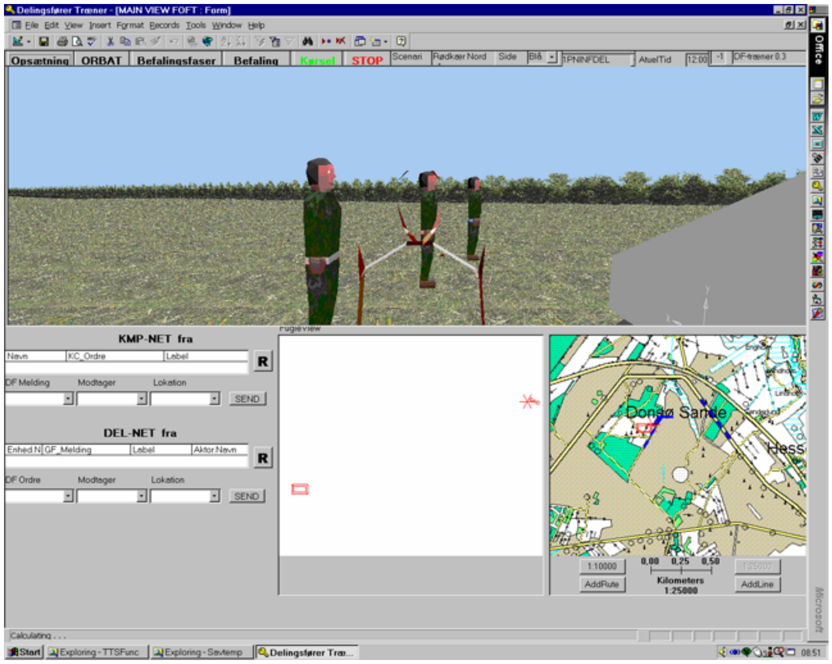

Figure D-35: Screen Dump from the Tactical Trainer for a Group Leader.

The development in Superscape has stopped, and the next generation of the VR will be done using Multigen Creator to create the "world" and DI-Guy for the soldiers.

D.3.3 F16 Flight Simulator

The F16 flight simulator is a readymade bought system from Hughes for the Air Force. It consists of a mock-up of a cockpit standing in front of 3 screens as shown in the picture. The simulation is run on a Silicon Graphic's computer from a control room.

Figure D-36: F16 Flight Simulator.

D.3.4 The Tank Simulator

The Tank simulator is a readymade bought system from Siemens (NL) and Simtech (Israel) for the Army. It is for shooting and battle exercises for platoon and below. The simulator has 4 40-foot containers. 3 of the containers have a technician room, a leopard 1A5 DK mock-up, a local instructor control panel, and a leopard 2A5 DK mock-up. The last container has a technician room, a central instructor control panel, and a report room with 12 seats. The set-up is shown in Figure D-37.

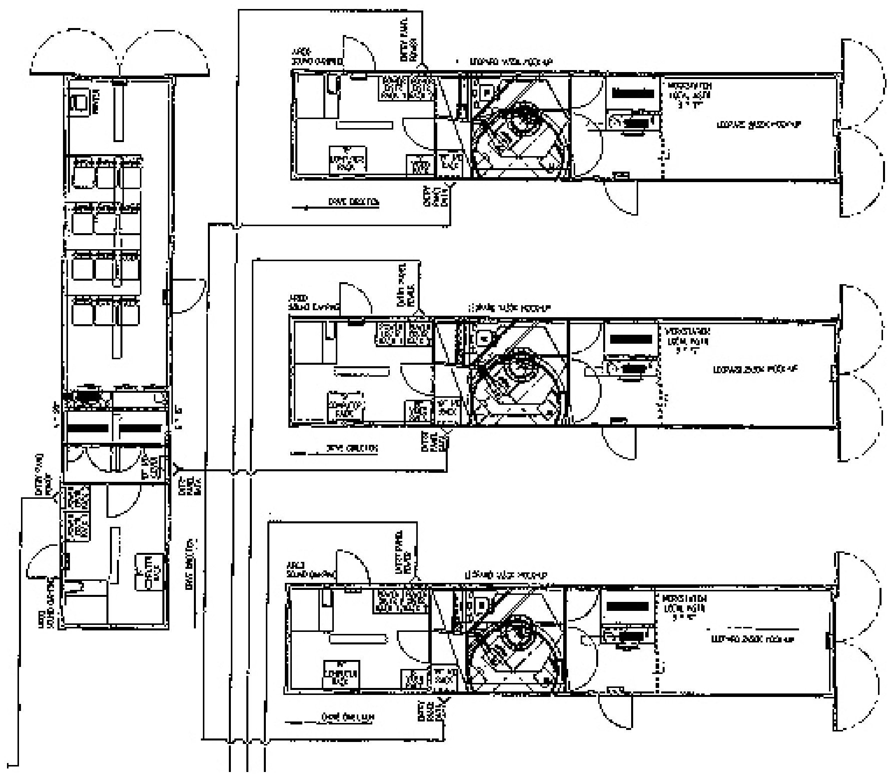

Figure D-37: Tank Simulator.

D.3.5 Finished Projects

D.3.5.1 Mission Debriefing System for the Royal Danish Air Force

Systematic (http://www.systematic.dk/) had developed a prototype for a debriefing system for aircraft crews to improve their skills based on experiences from completed missions. The system was developed on Commercial Off The Shelf (COTS) software and on a PC. It was a portable, low-cost Virtual Reality (VR) training system for aircraft crews. The benefit of the portability was that the system could be used anywhere the unit was deployed and by any crewmember.

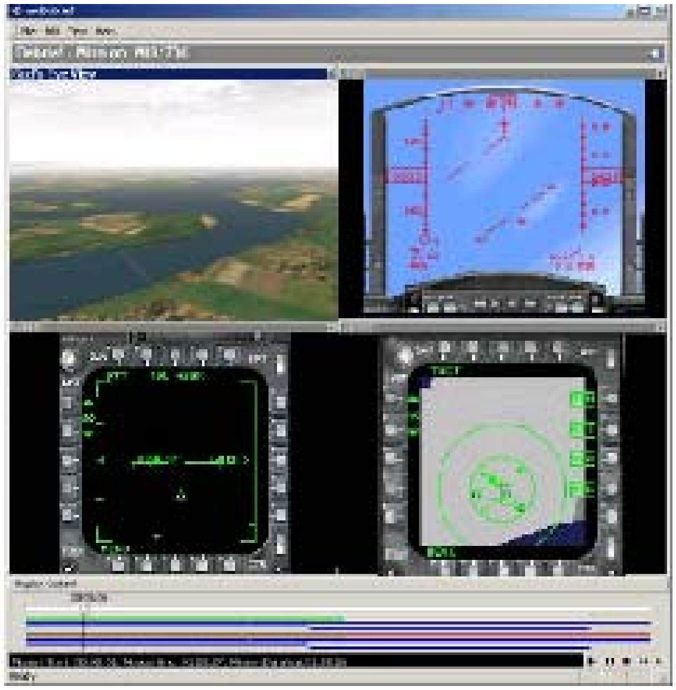

Figure D-38: Mission Debriefing System.

Flight hours are rather expensive and therefore the air forces must maximise the benefits from spent flight hours. This, combined with the fact that most air force units need to operate from different deployments remote from home bases, led the operational fighter squadrons to express a need for a low-cost debriefing system. The users were directly involved in the design and the focus was set on functionality – not technology. This approach resulted in a system, which gained acceptance among users and therefore became an everyday training tool. Driven by user requirements, the system was developed to run on a Microsoft Windows 2000 platform, and the system could interface with other systems. Furthermore, it had been essential to develop a system, which could be rapidly implemented.

The debriefing system used already existing information from the aircraft. The aircraft was equipped with Global Positioning System (GPS), three video cameras, and a microphone system to record the pilot's voice communication. The video cameras recorded the pilot's view through his head-up display and the entire instrument panel. Prior to the debriefing session all information from the aircraft (GPS-data, video- and audio recordings) was fed into the debriefing system. The GPS-data was loaded into a three dimensional (3D) model containing geographical information, the video and audio recordings were digitised, and all data were synchronised. On each monitor, four visual sources could be displayed concurrently, e.g. video recordings from three different aircraft and the graphical 3D view of the area, including aircraft. The selected visual sources were displayed along with a selected audio recording. The 3D graphic made it possible to see and follow selected aircraft from different perspectives on their mission. Furthermore, it was possible to see them chase other aircraft and to track their route by position, direction, and speed. The crew and other mission participants could by themselves prepare and execute the debriefing session.

The system was developed as a first generation version with basic functionality. For further information contact Major Birger I. Johansen, but the Royal Danish Air Force has lost interest in the system, so development has stopped.

D.3.5.2 RTP 6.14 Virtual Environment in Military Decision Making

The Western European Union (WEU) places contracts for research and development (R&D) projects under the European Co-operation for Long Term in Defence (EUCLID) programme. EUCLID's objectives are to provide European defence with basic technologies fully consistent with the aim of the Western European Armaments Organisation (WEAO) and to promote European industrial co-operation in R&D.

EUCLID consists of a set of Common European Priority Areas (CEPAs). CEPA 6 concerns "Advanced Information Processing and Communications". Within CEPA 6 the WEU is running a Research and Technology Project (RTP) number 6.14 entitled "Virtual Environment for military decision-making". The objectives of the RTP 6.14 project are to:

- Demonstrate the feasibility of using Virtual Reality (VR) techniques and a Virtual Environment (VE) in military decision-making in two functional areas: intelligence and logistics; and
- Show how the military decision-maker's situational awareness is improved by using a demonstrator developed in the RTP.

The scope of RTP 6.14 is limited to decision-making at battalion and brigade levels during wartime and peace support operations on the (battle)field.

The RTP 6.14 project is being performed by an industrial team, led by Atos Origin (The Netherlands). Atos Origin's partners are IFAD (Denmark), INTRACOM (Greece) and DATAMAT (Italy). The industrial team is named "Military Applications for VE in Logistics and intelligence" (MARVEL). The demonstrator to be developed during the MARVEL project will be named the "MARVEL demonstrator". The client is the Western European Armaments Organisation's Research Cell (WRC). A Management Group (MG) of the Ministries of Defence of Denmark, Greece, Italy, and The Netherlands manages the RTP 6.14 project on the WRC's behalf. The Dutch Ministry of Defence leads the MG.

There are four Work Packages (WPs) in the RTP 6.14 project:

- *WP 0 Management and Integration.* The objective of WP 0 is to manage the project and to integrate the various components into the MARVEL demonstrator.
- *WP 1 Preparation Elements.* The objective of WP 1 is to prepare for the research and implementation phases of the project.
- *WP 2 Research Elements.* The objective of WP 2 is to perform research in the various areas that are relevant for the evaluation of VE for military decision-making.
- *WP 3 Implementation and Evaluation of Demonstrator.* The objective of WP 3 is to implement and evaluate a demonstrator supporting military intelligence and logistics decision-making for the land-battle domain using VE display and user-input devices.
- *WP 4 Presentation.* The objective of WP 4 is to communicate the technical results of the project to parties outside the MARVEL consortium.

The project finished June 2003.

D.3.5.3 VIKING GIS

Denmark, Norway, and Sweden had a joint project. They wanted to design a common new submarine that could be purchased for all three nations. In this project different studies were made. One of these was

VIKING GIS. All the defence research establishments from the three countries participated in this study (which was the first common project between all three establishments).

The VIKING GIS is divided into two parts:

- VIKING GIS, Part I:
  - Establishing a database describing available data and its sources and quality.
  - Building a demonstrator to show how GIS can be used in the submarine.
- VIKING GIS, Part II.

D.3.5.3.1 VIKING GIS, Part I

The database will contain description of available digital and digitisable relevant marine geographic information from both civilian and military sources. An overview of future planed marine surveys shall be included. Special emphasis shall be given to the following parameters: depth, bottom characteristics, wrecks⁄objects, temperature and salinity or sound velocity profile, current, and acoustic and magnetic properties. The study shall include an assessment of data quality such as accuracy, resolution, and coverage. The database will have a graphical interface, where you can ask what information is available for defined areas.

A first version of a demonstrator package, named Concept Presenter (ConPres), shall be developed. Issues concerning modelling and visualisation of seabed terrain shall be addressed, especially evaluation of methods, technologies and functionality. ConPres shall demonstrate the use of GIS technologies for navigational purposes, more specifically:

- Use of several types of marine geographical information such as seabed topography, sediment types, acoustic profiles and wrecks/objects;
- Assess quality parameters as data resolution and data accuracy;
- Illustrate compression and visualisation techniques for large amounts of data; and
- Use of sensor information.

ConPres is composed of individual, more or less coupled components as articles, pictures and software.

D.3.5.3.2 VIKING GIS, Part II

Second part of Viking GIS focus on the advantages of GIS for the Viking submarine. This part is not a visualisation project, but it specifies how GIS and visualisation can be used in the future submarine.

The study comprises internal GI systems, and only deals with external systems to an extent necessary for the overall information flow. The main emphasis is on the role of the submarine as a naval combat platform with advanced sensors. The work does not comprise detailed technical solutions or demonstrators. The work includes discussions of technology with respect to feasibility.

The main objective of the study is to explore the utilization of advanced GIS technology in a submarine for more effective and secure realization and completion of today’s and tomorrow’s missions. Emphasis is put on the exploration of the emerging GIS capabilities with other IT advances to meet the challenges of the

21st century battle space. The study is carried out using different perspectives of how GIS can be used in the following selected missions:

- Reconnaissance;
- Special operation; and
- Mining and mine countermeasure.

D.4 CONCLUSION FOR DANISH VISUALISATION PROJECTS

The previous sections has mentioned several centres involved with visualisation of large datasets, and some of the visualisation projects going on in these centres. The table below is a summary. The hachure is centres, and the solid is projects. Red is military, and green is civilian. For each projects, the comments describe the type of project (divided into two categories), the techniques used, and if it is completed (if nothing is mentioned the project is ongoing).

Table D-2: Summary of Danish Projects

NAME	ESTABLISHMENT	COMMENTS
HCI-Lab	Danish Defence Research Establishment	Testing VR Software and Hardware
Tactical Trainer for a Group Leader	Danish Defence Research Establishment	R&D Simulation GIS, VR
F16 Flight Simulator	The Air Force	Commercial Simulation GIS, VR Completed
Tank Simulator	The Army	Commercial Simulation Completed
Debriefing System for the Air Force	Systematic	Commercial Simulation 3D Disconnected
RTP 6.14 Virtual Environment in Military Decision Making	Atos Origin (The Netherlands), IFAD (Denmark), INTRACOM (Greece), and DATAMAT (Italy)	Commercial, R&D Decision Support GIS, VR Concluded
VIKING GIS	Danish Defence Research Establishment, Norwegian Defence Research Establishment, Swedish Defence Research Agency	R&D (Study) GIS, VR 3D Concluded
Centre for Pervasive Computing		Most of the work in the centre is organised as Research Projects involving both companies and universities

NAME	ESTABLISHMENT	COMMENTS
Centre for Advanced Visualisation and Interaction – CAVI	Centre for Pervasive Computing	3D Panorama Cinema TAN Holobench Virtual Studio
VR Media Lab	University of Aalborg	Cave with 6 Walls Powerwall (large flat screen) Panoramic Screen (160° curved large screen)
Centre for 3DGI	European Regional Development Fund (ERDF), Aalborg University, the National Survey and Cadastre – Denmark, Kampsax A/S, and Informi GIS A/S	VR Media Lab
VR•C	The Technical University, UNI-C	TAN Powerwall (large flat screen) 3 n-vision DataVisor HiRes Head Mounted Display See-Through (augmented reality)
FORCE Technology	FORCE Technology	Mock-up of Bridges on Ships
Virtual Urban Planning	Aarhus University, Cadpeople, and COWI	Commercial, R&D Landscape Decision Support GIS, VR
Visualisation of the Cadiovascular System	CAVI, Aarhus University Hospital, and Systematic Software Engineering	R&D Medical VR
Digital Theatre – Hyperopticon	Aarhus University (Denmark), the Academy of Figurative Theatre (Norway), Studio di Progettazione (Italy), and CAVI (Denmark)	Commercial, R&D Entertainment VR
Video Prototyping	Aarhus University	R&D Medical Video
3D Experiences	Aarhus University and Zentropa Interaction	Commercial, R&D Entertainment and Education 3D Interactive TV
3D Sound in 3D Space	Aarhus University and TC Electronic	R&D 3D Sound in 3D Models
Interfaces for 3D Applications	Aarhus University, Personics Aps, Systematic Software Engineering A/S, RoninWorks BV (Holland) and CAVI	Commercial, R&D VR 3D
Sound as Media, between Signal and Music	DAIMI, Aesthetic Disciplines, IMV, Engineering College of Aarhus, TC Electronic (R&D), DIEM, The Music Conservatory (Aarhus), LEGO (R&D), and Mosart network	Commercial, R&D Sound

NAME	ESTABLISHMENT	COMMENTS
Distributed Multimedia Technologies and Applications (DMM)	Intermedia, Aarhus University	R&D Hypermedia VR Real-Time Multi-Modal Communication
3D Image Processing for Cranium- and Brain-Surgical Planning and Simulation	CAVI, PET Centre	R&D Medical 3D Concluded
Digital, 3D Atlas of the Receptor Systems of the Human Brain	CAVI, PET Centre	R&D Medical 3D Concluded
Packaging Design	CAVI, Technical University of Copenhagen, Arla Foods	Commercial 3D Concluded
Landscape Visualisation	CAVI, Technical University of Copenhagen, Kampsax/Lundbeck	Commercial, R&D Landscape 3D Concluded
Architecture	CAVI, Technical University of Copenhagen, Arkitektfirmaet Schütze A/S	Commercial, R&D Architecture 3D Concluded
City Planning	CAVI, Technical University of Copenhagen, CF Møller/CADpeople	Commercial, R&D Architecture Landscape 3D Concluded
Virtual Set Activities / Ren Kagemand (Danish Television Show)	CAVI, Technical University of Copenhagen, Danmarks Radio	Commercial Entertainment VR Concluded
Beaching Museum St. George	CAVI, Technical University of Copenhagen, Jydsk Dykkerfirma (Diving firm in Jutland)	Commercial Architecture 3D Concluded
Family Factory	CAVI, the Danish Film School, Schule für Schauspiel-kunst “Ernst Bush” Berlin	Commercial, R&D Entertainment 3D Concluded
Digital Theatre	CAVI	R&D Entertainment 3D Concluded

NAME	ESTABLISHMENT	COMMENTS
Whizbang	CAVI	R&D 3D Concluded
Aarhus New Art Museum	CAVI, the art department Schmidt, Hammer and Lassen	Commercial Architecture 3D Concluded
CAVI Building	CAVI	Commercial Architecture 3D Concluded
Katrinebjerg Phase 1	CAVI	Commercial Architecture Landscape 3D Concluded
Katrinebjerg Phase 2	CAVI	Commercial Architecture Landscape 3D Concluded
Architectural Competition	CAVI, the architecture firms of Jørn Schütze	Commercial Architecture 3D Concluded
Project of Sonderborg	VR Media Lab, Cadpeople and COWI	Commercial Landscape VR
Aalborg University Campus Model	VR Media Lab	R&D Architecture VR
CAE and CFD into Virtual Reality (CCVR)	VR Media Lab	R&D Scientific VR
3D Airflow	VR Media Lab, the Department of Building Technology and Structural Engineering (also at Aalborg University), the Danish Institute of Agricultural Sciences, and the Royal Veterinary and Agricultural University	R&D Environmental Scientific 3D
Data Mining	VR Media Lab	R&D Education Scientific 3D VR

NAME	ESTABLISHMENT	COMMENTS
3D Visual Data Mining	VR Media Lab	R&D Education Scientific 3D VR
Learning	VR Media Lab, Vestas and Zenaria	Commercial, R&D Education VR
Vestas	VR Media Lab	Commercial Education Architecture VR
Music House in Aalborg	VR Media Lab and the Danish National Research and Educational Buildings	Commercial Architecture VR Concluded
A Plan for New Housing Development	VR Media Lab, COWI	Commercial Architecture Landscape VR Concluded
Highway at Holbæk	VR Media Lab	Commercial Landscape VR Concluded
Model of NOVI 3, 4, and 5	VR Media Lab	Commercial Architecture VR Concluded
Puppet	University of Aalborg	R&D Education VR Concluded
Staging of Virtual 3D-spaces	University of Aalborg, etc.	Commercial, R&D 3D VR Multimedia Concluded
Sound in Cyberspace	University of Aalborg	R&D Sound Concluded
Qualification of 3D GeoInformation	Centre for 3DGI	Commercial, R&D Landscape GIS, VR 3D

NAME	ESTABLISHMENT	COMMENTS
Queries for Navigation in 3D Models	Centre for 3DGI	Commercial, R&D Landscape GIS, VR 3D
3DGI Viewer	Centre for 3DGI	Commercial, R&D Landscape GIS, VR 3D
Interactive 3D Visualisation of Projects	VR•C	R&D 3D
Scientific Visualisation	VR•C	R&D Scientific
Network-Based VR Technology	VR•C	R&D VR Concluded
VR in Neuro Informatic	The Technical University	R&D Medical VR Concluded
Simulating a Combine Harvester	The Technical University	R&D VR Simulation Concluded
Ship Simulation	FORCE Technology	Commercial Simulation GIS, VR
Skyline	COWI	Commercial Landscape GIS 3D
Vision Quest	UNI-C	
Digital 3D-Architecture Models on the Internet	Art Academy Architecture School	Commercial, R&D 3D Concluded
Visualising the Information in a Library System	School of Commerce	Commercial, R&D (Prototype) Concluded
Motion Capture	3D Connection	Commercial VR 3D

	= Military
	= Civilian
	= Centre

As the table shows most projects can be categorised in the following categories:

- Landscape;
- Architecture;
- Decision Support;
- Simulation;
- Scientific;
- Medical;
- Entertainment; and
- Education.

The categories are mentioned in the order of interest for the military.

The visualisation techniques are:

- 3D;
- VR;
- GIS;
- Sound;
- 3D sound;
- Multimedia;
- Hypermedia;
- Real-time multi-modal communication;
- Interactive TV; and
- Video.

Visualisation and VR are fertile research areas in Denmark. The interaction between the military research and the civilian universities and research centres is improving. Several civilian research projects show great potential for military use.

Annex E – NORWEGIAN WORK ON VISUALISATION

Tore Smestad
Norwegian Defence Research Laboratory (FFI)

E.1 INTRODUCTION

This chapter contains the most recent of a series of annual updates on national work on visualisation for Norway, written by the Norwegian Delegate. These elements are hopefully representative of the visualisation research and activity in Norway, but by no means complete. The elements of each reported items are:

- **Item:** A title of the item containing some data for enabling identification.
- **Category:** Grouping the type of report elements (here only textually described).
- **Description:** Description as long or as short as the presenter deems necessary.
- **References:** To make a reader able to get hold of additional information, e.g. web-address.
- **Update:** Who last put in or checked the data and when.

Item: **Major research and education units in Norway involved in visualisation**

Category: **Summary**

Description: This list hopefully includes most of such units having current and noticeably activities in visualisation. The references are made to a general web page, many of which are written in Norwegian. Hopefully, the information found here might be to some help for further search and references. Most of the other reporting elements of this update are result of activities in some of these institutions. The institutions are grouped alphabetically under their location in Norway.

E.2 LISTING OF SITES

Oslo: FFI (www.ffi.no), SINTEF in Oslo (www.oslo.sintef.no).

Trondheim: HIST (www.hist.no), NTNU (www.ntnu.no) especially at ITK (www.itk.ntnu.no).

Bergen: CMR (www.cmr.no).

Halden: (Site of NATO Workshop Sept 10-13 2002) IFE (www.ife.no).

E.3 LISTING OF PROJECTS

E.3.1 CMR Bergen

Item: **A VR-laboratory for multiple use**

Category: **Application of VR**

Description: The research institute CMR in Bergen has applied VR-technology in different area, especially aspects of oil-production. They have a VR laboratory to demonstrate the potential in the field of medicine, architecture, regional planning and climate studies.

References: See under "Departments: Advanced computing/Facilities" under www.cmr.no.

E.3.2 Viz at FFI

Item: **Transparent volume visualisation by "Viz" used for flow phenomena in 3D**

Category: **Software and application experience**

Description: Numeric simulations of flow phenomena in 3D at FFI called for advanced 3D visualisation. Software was then developed for powerful Silicon Graphics machines. The software now called "Viz" is further developed at FFI and made more general. Viz is made for extremely fast voxel based volume rendering. Many groups working in the field of high performance computing use Viz for visualisation (see ftp://ftp.ffi.no/spub/stsk/viz/index.html). Most significant is perhaps the use Viz by some of the ASCI (Accelerated Strategic Computing Initiative) groups in USA.

J.O. Langseth at FFI produced an illustration for the Springer Calendar MathInsight 2002 (November) by this type of visualisation (see http://www.math.uiuc.edu/~gfrancis/public/13uebersicht.pdf). Anders Helgeland and Øyvind Andreassen from FFI won the award for best poster at the IEEE Visualisation conference in Boston 2002 (see http://vis.computer.org/vis2002/special/paper-award.html).

References: See and J.O. Langseth: "3D visualisation of shock waves using volume rendering" to appear in Proc. Godunov Methods: Theory and Applications, Oxford, 1999.

E.3.3 Commercial Application of VR

Item: **Use of VR in petroleum production offshore in Norway**

Category: **Application experience**

Description: Research activities in the institute CMR in Bergen resulted in a firm applying VR-technology in different aspects of petroleum production, now obviously bought by Schlumberger. The Norwegian firm "Norsk Hydro" heavily engaged in offshore petroleum production, has played a big role here.

References: http://www.sis.slb.com/content/about/forum/2002/abstract_mdrd_012.asp and www.gasandoil.com/goc/company/cne95016.htm

Update: Tore Smestad, September 2003.

E.3.4 Statoil Change and Learning Lab

Item: **HIST educates students using Statoil facilities for visualisation of petroleum production offshore in Norway**

Category: **Education and applications**

Description: A cooperation started in the autumn 1999 between HIST (education near the University level) and the oil company "Statoil" in Trondheim for visualisation of work operations in offshore oil production. It has resulted in the so called "Statoil′s Change and Learning Lab" (SELL). Approximately 20 students each year choose this possibility.

References: A start address is http://129.241.149.139/ (most in Norwegian). See also www.gasandoil.com/goc/company/cne95016.htm

E.3.5 Nuclear Safety Laboratory in Halden

Item: **The HAMMLAB and the Virtual Reality Centre in Halden**

Category: **Laboratory facilities**

Description: HAMMLAB is a laboratory facility for experiments in control room situations motivated by international cooperation in nuclear power plant safety (the OECD Halden reactor project). Central topics have been situation awareness and development of software to assist humans.

References: See www.ife.no/english/tillegg/index.jsp?avdelingsId=1963&tilleggId=1228 for HAMMLAB and for the OECD Halden reactor project: http://www.ife.no/english/avdelinger/avdeling_details.jsp?avdelingsId=1963&nav=1, with a link to "Halden Virtual Reality Centre". One example of application: http://www.ife.no/media/360_VR-reaktor.pdf

E.3.6 ITK at NTNU Trondheim

Item: **"KIKS" – a multidisciplinary program on body-carried information – and communication systems**

Category: **Research program**

Description: The vision for a carried information- and communication system is that the operation in a process plant will have access to all necessary information at all time end places, such that this type of equipment may substitute today's central control rooms.

References: http://www.itk.ntnu.no/KIKS/kiks.htm (in Norwegian).

E.3.7 Det Norske Veritas

Item: **Tests and research of wearable computers in maritime application**

Category: **Research activity**

Description: The maritime classification company "Det Norske Veritas" (DNV) was involved in testing a prototype of a wearable computer (from IBM) from December 1999 to March 2000 for ship inspection. This is an example of research activity on wearable computers for maritime applications.

References: A description of this was found in a Norwegian technical magazine ("Teknisk Ukeblad" no 16 April 2001, written in Norwegian (link no more available). An other reference to this type of general activity in Norway is a paper found at http://www.dcs.gla.ac.uk/~johnson/papers/mobile/HCIMD1.html titled "Developing Scenarios for Mobile CSCW". Further references might be found in the paper. See also the item above about "KIKS").

E.3.8 Octaga

Item: **Norwegian firm Octaga**

Category: **Business, development, products**

Description: The firm Octaga was established in 2001 based on research activity on 3D and VR in Telenor R&D; Telenor is the major telecom and internet operator in Norway. They describe the firm as multidisciplinary with own developed competitive software to be run on PC's. Examples of projects are found on their homepage.

References: Info on homepage: www.octaga.com.

E.3.9 Norwegian Armed Forces

Item: **Simulation systems and VR used in the Norwegian Armed Forces**

Category: **Military education and training**

Description: Installations for military educations and training using different VR technology in simulation systems are found in all three services in Norway. The main purpose is to make the education and training more cost effective. The systems are used for a number of different weapons and weapon systems, both for individual training and for different sizes of teams. Most of the systems for the Army are found at Rena while those for the Navy are located in Bergen. The Air Force has installed systems at Rygge and Bodø.

References: A NATO RTO report describing member Nations state of VR in military research and applications can be found at http://www.rta.nato.int/Rdp.asp?RDP=RTO-TR-018.

E.3.10 Norwegian Battle Lab Experimentation in Bodø

Item: **"Common Operation Decision System" – CODS**

Category: **Visualisation device**

Description: This system applies high resolution LCD screens configured 3x3, 4x4 or 5x5 to form a very high resolution display surface to present basically geo-related data for commanders and decision teams in military operations. It allows "seamless" transition from strategic to tactical level and all scales all over the world given maps/pictures, etc., presented in "2½ D". The system is developed in the ongoing "Norwegian Battle Lab Experimentation" (NOBLE) in Bodø. It is installed for experimentation and evaluation in different command centers in Norway, also used in Kirghizstan in relation to Norwegian Air Force participation in operation "Enduring Freedom". It was presented at the NATO Workshop IST-036/RWS-005 "Massive Military Data Fusion and Visualisation – Users Talk with Developers" in Halden Norway, September 2002.

References: About NOBLE: http://www.mil.no/prosjekter/noble/noble_english/start/NOBLE/

Homepage of a firm involved (look for ADVICORE): http://www.hep.no.

Documentation from IST-036/RWS-005 can be found on http://www.rta.nato.int/

E.3.11 Norwegian Defence Research Lab

Item: **Command and Control/Combat Management Lab at FFI**

Category: **Laboratory facilities able for visualisation research**

Description: This laboratory was built in connection with a now completed project at the Norwegian Defence Research Establishment (FFI), and might be used for different purposes included data fusion and visualisation experiments/research. The lab has a heterogeneous computer network with workstations and PCs. It runs software for visualisation and for simulating different platforms, missiles, sensors, and communication systems, and is based on High Level Architecture (HLA). The lab is equipped with a 3D table using passive stereo, projectors for presenting multiple large screen displays, two 3-screen panoramic display units, in addition to the screens associated with the different workstations and PCs. The use of COTS in Combat Management Systems has been an important focus up to now – see Figure E-1.

References: Contact person: Chief Scientist Karsten Bråthen, e-mail: kab@ffi.no.

Figure E-1: Command and Control / Combat Management Lab at FFI.

REPORT DOCUMENTATION PAGE

1. Recipient's Reference	2. Originator's References	3. Further Reference	4. Security Classification of Document
	RTO-TR-IST-021 AC/323(IST-021)TP/43	ISBN 978-92-837-0067-8	UNCLASSIFIED/ UNLIMITED

5. Originator Research and Technology Organisation
North Atlantic Treaty Organisation
BP 25, F-92201 Neuilly-sur-Seine Cedex, France

6. Title
Multimedia Visualisation of Massive Military Datasets

7. Presented at/Sponsored by
This Report documents the findings of Task Group IST-021/RTG-007.

8. Author(s)/Editor(s)	9. Date
Multiple	June 2007

10. Author's/Editor's Address	11. Pages
Multiple	252

12. Distribution Statement There are no restrictions on the distribution of this document. Information about the availability of this and other RTO unclassified publications is given on the back cover.

13. Keywords/Descriptors

Communications networks	Information systems	Specifications
Computer networks	Integrated systems	Standards
Correlation techniques	Models	Systems analysis
Data fusion	Network visualisation	Systems engineering
Data visualisation	Software development	Visualisation techniques
Human factors engineering	Software engineering	Visualisation technologies
Information operations		

14. Abstract

IST-021/RTG-007 considered the visualisation of massive military datasets, with a particular interest in methods of evaluating visualisation systems. In common with its predecessor groups under both RTO and the earlier Defence Research Group, IST-021/RTG-007 considered visualisation to happen inside a person's head, not on a screen. For evaluation purposes, several different nations offered national visualisation systems, but for different individual reasons only two were finally made available. Two evaluation techniques were used, both of which presented some difficulties for evaluators when confronted with the complexities of real military visualisation systems. A "visualisation testbed" was constructed and used for several studies of display techniques in support of visualisation. Under the patronage of IST-021/RTG-007, a visualisation network of experts (NX) was continued, and both the NX and IST-021/RTG-007 organized workshops that examined particular facets of the visualisation problem, with some emphasis on counter-terrorism operations. The work of the group demonstrated that a major visualisation issue was the representation of network phenomena, such as computer network security against attack, social networks in counter-terrorism, causal and supportive networks in logistical operations, networks of effects in infrastructure damage, and a wide variety of other network situations of military relevance. The group recommended that a technical team be set up to study the visualisation of network-related phenomena. Such a group was approved with the designation IST-059/RTG-025.

NORTH ATLANTIC TREATY ORGANISATION

RESEARCH AND TECHNOLOGY ORGANISATION

BP 25
F-92201 NEUILLY-SUR-SEINE CEDEX • FRANCE
Télécopie 0(1)55.61.22.99 • E-mail mailbox@rta.nato.int

DIFFUSION DES PUBLICATIONS RTO NON CLASSIFIEES

Les publications de l'AGARD et de la RTO peuvent parfois être obtenues auprès des centres nationaux de distribution indiqués ci-dessous. Si vous souhaitez recevoir toutes les publications de la RTO, ou simplement celles qui concernent certains Panels, vous pouvez demander d'être inclus soit à titre personnel, soit au nom de votre organisation, sur la liste d'envoi.

Les publications de la RTO et de l'AGARD sont également en vente auprès des agences de vente indiquées ci-dessous.

Les demandes de documents RTO ou AGARD doivent comporter la dénomination « RTO » ou « AGARD » selon le cas, suivi du numéro de série. Des informations analogues, telles que le titre est la date de publication sont souhaitables.

Si vous souhaitez recevoir une notification électronique de la disponibilité des rapports de la RTO au fur et à mesure de leur publication, vous pouvez consulter notre site Web (www.rta.nato.int) et vous abonner à ce service.

CENTRES DE DIFFUSION NATIONAUX

ALLEMAGNE
Streitkräfteamt / Abteilung III
Fachinformationszentrum der Bundeswehr (FIZBw)
Gorch-Fock-Straße 7, D-53229 Bonn

BELGIQUE
Royal High Institute for Defence – KHID/IRSD/RHID
Management of Scientific & Technological Research for Defence, National RTO Coordinator
Royal Military Academy – Campus Renaissance
Renaissancelaan 30
1000 Bruxelles

CANADA
DSIGRD2 – Bibliothécaire des ressources du savoir
R et D pour la défense Canada
Ministère de la Défense nationale
305, rue Rideau, 9e étage
Ottawa, Ontario K1A 0K2

DANEMARK
Danish Acquisition and Logistics Organization (DALO)
Lautrupbjerg 1-5
2750 Ballerup

ESPAGNE
SDG TECEN / DGAM
C/ Arturo Soria 289
Madrid 28033

ETATS-UNIS
NASA Center for AeroSpace Information (CASI)
7115 Standard Drive
Hanover, MD 21076-1320

FRANCE
O.N.E.R.A. (ISP)
29, Avenue de la Division Leclerc
BP 72, 92322 Châtillon Cedex

GRECE (Correspondant)
Defence Industry & Research General Directorate, Research Directorate
Fakinos Base Camp, S.T.G. 1020
Holargos, Athens

HONGRIE
Department for Scientific Analysis
Institute of Military Technology
Ministry of Defence
P O Box 26
H-1525 Budapest

ISLANDE
Director of Aviation
c/o Flugrad
Reykjavik

ITALIE
General Secretariat of Defence and National Armaments Directorate
5th Department – Technological Research
Via XX Settembre 123
00187 Roma

LUXEMBOURG
Voir Belgique

NORVEGE
Norwegian Defence Research Establishment
Attn: Biblioteket
P.O. Box 25
NO-2007 Kjeller

PAYS-BAS
Royal Netherlands Military Academy Library
P.O. Box 90.002
4800 PA Breda

POLOGNE
Centralny Ośrodek Naukowej Informacji Wojskowej
Al. Jerozolimskie 97
00-909 Warszawa

PORTUGAL
Estado Maior da Força Aérea
SDFA – Centro de Documentação
Alfragide
P-2720 Amadora

REPUBLIQUE TCHEQUE
LOM PRAHA s. p.
o. z. VTÚLaPVO
Mladoboleslavská 944
PO Box 18
197 21 Praha 9

ROUMANIE
Romanian National Distribution Centre
Armaments Department
9-11, Drumul Taberei Street
Sector 6
061353, Bucharest

ROYAUME-UNI
Dstl Knowledge Services
Information Centre
Building 247
Dstl Porton Down
Salisbury
Wiltshire SP4 0JQ

TURQUIE
Milli Savunma Bakanlığı (MSB)
ARGE ve Teknoloji Dairesi Başkanlığı
06650 Bakanliklar
Ankara

AGENCES DE VENTE

NASA Center for AeroSpace Information (CASI)
7115 Standard Drive
Hanover, MD 21076-1320
ETATS-UNIS

The British Library Document Supply Centre
Boston Spa, Wetherby
West Yorkshire LS23 7BQ
ROYAUME-UNI

Canada Institute for Scientific and Technical Information (CISTI)
National Research Council Acquisitions
Montreal Road, Building M-55
Ottawa K1A 0S2, CANADA

Les demandes de documents RTO ou AGARD doivent comporter la dénomination « RTO » ou « AGARD » selon le cas, suivie du numéro de série (par exemple AGARD-AG-315). Des informations analogues, telles que le titre et la date de publication sont souhaitables. Des références bibliographiques complètes ainsi que des résumés des publications RTO et AGARD figurent dans les journaux suivants :

Scientific and Technical Aerospace Reports (STAR)
STAR peut être consulté en ligne au localisateur de ressources uniformes (URL) suivant:
http://www.sti.nasa.gov/Pubs/star/Star.html
STAR est édité par CASI dans le cadre du programme NASA d'information scientifique et technique (STI)
STI Program Office, MS 157A
NASA Langley Research Center
Hampton, Virginia 23681-0001
ETATS-UNIS

Government Reports Announcements & Index (GRA&I)
publié par le National Technical Information Service
Springfield
Virginia 2216
ETATS-UNIS
(accessible également en mode interactif dans la base de données bibliographiques en ligne du NTIS, et sur CD-ROM)

NORTH ATLANTIC TREATY ORGANISATION

RESEARCH AND TECHNOLOGY ORGANISATION

BP 25

F-92201 NEUILLY-SUR-SEINE CEDEX • FRANCE

Télécopie 0(1)55.61.22.99 • E-mail mailbox@rta.nato.int

DISTRIBUTION OF UNCLASSIFIED RTO PUBLICATIONS

AGARD & RTO publications are sometimes available from the National Distribution Centres listed below. If you wish to receive all RTO reports, or just those relating to one or more specific RTO Panels, they may be willing to include you (or your Organisation) in their distribution.

RTO and AGARD reports may also be purchased from the Sales Agencies listed below.

Requests for RTO or AGARD documents should include the word 'RTO' or 'AGARD', as appropriate, followed by the serial number. Collateral information such as title and publication date is desirable.

If you wish to receive electronic notification of RTO reports as they are published, please visit our website (www.rta.nato.int) from where you can register for this service.

NATIONAL DISTRIBUTION CENTRES

BELGIUM
Royal High Institute for Defence – KHID/IRSD/RHID
Management of Scientific & Technological Research for Defence, National RTO Coordinator
Royal Military Academy – Campus Renaissance
Renaissancelaan 30
1000 Brussels

CANADA
DRDKIM2
Knowledge Resources Librarian
Defence R&D Canada
Department of National Defence
305 Rideau Street, 9th Floor
Ottawa, Ontario K1A 0K2

CZECH REPUBLIC
LOM PRAHA s. p.
o. z. VTÚLaPVO
Mladoboleslavská 944
PO Box 18
197 21 Praha 9

DENMARK
Danish Acquisition and Logistics Organization (DALO)
Lautrupbjerg 1-5
2750 Ballerup

FRANCE
O.N.E.R.A. (ISP)
29, Avenue de la Division Leclerc
BP 72
92322 Châtillon Cedex

GERMANY
Streitkräfteamt / Abteilung III
Fachinformationszentrum der Bundeswehr (FIZBw)
Gorch-Fock-Straße 7
D-53229 Bonn

GREECE (Point of Contact)
Defence Industry & Research General Directorate
Research Directorate
Fakinos Base Camp
S.T.G. 1020
Holargos, Athens

HUNGARY
Department for Scientific Analysis
Institute of Military Technology
Ministry of Defence
P O Box 26
H-1525 Budapest

ICELAND
Director of Aviation
c/o Flugrad, Reykjavik

ITALY
General Secretariat of Defence and National Armaments Directorate
5th Department – Technological Research
Via XX Settembre 123
00187 Roma

LUXEMBOURG
See Belgium

NETHERLANDS
Royal Netherlands Military Academy Library
P.O. Box 90.002
4800 PA Breda

NORWAY
Norwegian Defence Research Establishment, Attn: Biblioteket
P.O. Box 25
NO-2007 Kjeller

POLAND
Centralny Ośrodek Naukowej Informacji Wojskowej
Al. Jerozolimskie 97
00-909 Warszawa

PORTUGAL
Estado Maior da Força Aérea
SDFA – Centro de Documentação
Alfragide
P-2720 Amadora

ROMANIA
Romanian National Distribution Centre
Armaments Department
9-11, Drumul Taberei Street
Sector 6, 061353, Bucharest

SPAIN
SDG TECEN / DGAM
C/ Arturo Soria 289
Madrid 28033

TURKEY
Milli Savunma Bakanlığı (MSB)
ARGE ve Teknoloji Dairesi Başkanlığı
06650 Bakanliklar – Ankara

UNITED KINGDOM
Dstl Knowledge Services
Information Centre
Building 247, Dstl Porton Down
Salisbury, Wiltshire SP4 0JQ

UNITED STATES
NASA Center for AeroSpace Information (CASI)
7115 Standard Drive
Hanover, MD 21076-1320

SALES AGENCIES

NASA Center for AeroSpace Information (CASI)
7115 Standard Drive
Hanover, MD 21076-1320
UNITED STATES

The British Library Document Supply Centre
Boston Spa, Wetherby
West Yorkshire LS23 7BQ
UNITED KINGDOM

Canada Institute for Scientific and Technical Information (CISTI)
National Research Council Acquisitions
Montreal Road, Building M-55
Ottawa K1A 0S2, CANADA

Requests for RTO or AGARD documents should include the word 'RTO' or 'AGARD', as appropriate, followed by the serial number (for example AGARD-AG-315). Collateral information such as title and publication date is desirable. Full bibliographical references and abstracts of RTO and AGARD publications are given in the following journals:

Scientific and Technical Aerospace Reports (STAR)
STAR is available on-line at the following uniform resource locator:
http://www.sti.nasa.gov/Pubs/star/Star.html
STAR is published by CASI for the NASA Scientific and Technical Information (STI) Program
STI Program Office, MS 157A
NASA Langley Research Center
Hampton, Virginia 23681-0001
UNITED STATES

Government Reports Announcements & Index (GRA&I)
published by the National Technical Information Service
Springfield
Virginia 2216
UNITED STATES
(also available online in the NTIS Bibliographic Database or on CD-ROM)

ISBN 978-92-837-0067-8