A Risk-Based Approach for Improving Intelligence, Surveillance, and Reconnaissance (ISR) Capabilities

Yvan Gauthier[1,2], **Neil Carson**[1,3], **Sean Bourdon**[1,4]
[1]Defence Research and Development Canada – Centre for Operational Research and Analysis
[2]Canada Command Operational Research and Analysis Team
[3]NORAD Operational Research Team
[4]RCAF Directorate of Air Staff Operational Research
National Defence Headquarters, 101 Colonel By Drive
Ottawa, ON, K1A 0K2

yvan.gauthier@drdc-rddc.gc.ca, neil.carson@drdc-rddc.gc.ca, sean.bourdon@drdc-rddc.gc.ca

ABSTRACT

Defence and security organizations are continuously trying to improve how they develop and employ intelligence, surveillance, and reconnaissance (ISR) capabilities. Strategic and operational decisions are often based on metrics that focus exclusively on vulnerability, such as detection, identification, or tracking metrics. Such metrics may adequately describe the performance and effectiveness of ISR capabilities in the context of specific scenarios. However, they do not necessarily provide means of directly comparing how well these capabilities fare against a broad range of threats. These metrics also tend to ignore the various consequences that threats may have if authorities fail to react in a timely and appropriate fashion. This paper proposes a quantitative, all-hazards risk assessment model for comparing different capability options. It decomposes the triad of threat, vulnerability and consequence into lower-level risk factors that can be assessed from various sources, such as intelligence, historical analysis, operational performance data, modelling, or expert judgement. Uncertainties surrounding the various risk factors are handled using fuzzy sets. Once assessed, risk factors are combined through fuzzy logic and fault-tree analysis in order to generate a risk profile, which forms the basis of an objective function for optimizing ISR capabilities. An example of how this approach can be applied to improving maritime domain awareness in a domestic security context is presented. Applications beyond the maritime domain across multiple environments and scenarios are also possible.

1.0 INTRODUCTION

Defence and security organizations are continuously trying to improve how they develop and employ intelligence, surveillance, and reconnaissance (ISR) capabilities. At a tactical level, several techniques exist to analyze and optimize these capabilities in the context of specific scenarios. At operational and strategic levels however, the process of defining what capability mix is required and how it should be employed is rarely optimized. The need for "persistent ISR", often called for in policy documents, is not well defined and generally unachievable over very large areas. More specific ISR requirements are needed to direct capability development, but setting and prioritizing these requirements in a rigorous fashion remains challenging, for multiple reasons.

1.1 Issues

1.1.1 Multiplicity of threats, hazards, mandates, and jurisdictions

One of the difficulties is that government departments and agencies must be able to detect, monitor, and respond to a wide spectrum of threats and hazards, from defence issues (e.g., military incursions) to security issues (e.g., terrorism, smuggling) and public safety issues (e.g., natural disasters, pandemics). Individual organizations within a government – or even within a same department – will prioritize ISR assets and capabilities that contribute the most to the achievement of their own mandates, according to their own metrics. Naturally, these tend to focus on the specific objectives of each organization and may not be directly comparable. But from a strategic perspective, since ISR assets are often shared by multiple organizations, such a comparison should be made in order to prioritize ISR requirements and potential solutions from a whole-of-government perspective.

1.1.2 Multiplicity of environments

An additional dimension of the problem is that multiple environmental domains (*maritime*, *land*, *air*, *space*, *human*, and *cyber* domains) must be monitored simultaneously. Most ISR assets can be used across multiple environments, an example of which is a patrol aircraft that may generate air, land, and maritime ISR products. However, when different environments are analyzed in isolation, the ISR solutions identified may not be as optimal or cost-effective as they could be, since the cross-domain effects of certain assets may not be recognized or taken into account.

1.1.3 Complexity of ISR systems and architectures

Another complicating factor is the variety of ISR sensors, systems, and sources that are used for developing and maintaining situational awareness, and their complex relationships. In part because of this complexity, many ISR studies and metrics focus on the individual *performance* of ISR assets, that is, their ability to accurately report target data (e.g., location, size, movement, signature, type, activity, disposition, identity) or other situational or environmental data. In addition, a significant focus is placed on data fusion and the generation of common operational pictures for representing the situational information. Relatively fewer studies concentrate on the *effectiveness* of ISR systems and architectures, that is, their ability to enable the achievement of particular force's or government's objectives. This is largely due to the difficulty in analyzing these complex systems in a way that is rigorous, transparent, and timely enough to influence strategic decisions.

1.1.4 Isolation of the *Sense* functions

ISR capabilities do not only influence *Sense* functions related to the acquisition and processing of information. They are critical enablers for many capabilities in other functional domains[1] such as the *Command*, *Shield*, or *Act* domains. For example, the ISR requirements associated with a specific scenario should take into account the extent to which potential targets need to be protected (*Shield* domain), since high-value targets may demand more surveillance. The requirements should also take into account how and where potential threats should be neutralized (*Act* domain). By considering, to some level, the relationships between functional domains, it becomes possible to identify ISR solutions that maximize overall force effectiveness. As shown in Table 1, force effectiveness metrics essentially focus on outcomes and, especially in a domestic context, the *risks* posed by

[1]The Canadian Forces (CF) use functional domains to categorize the various capabilities they require for routine and contingency operations. The domains are defined in the CF Integrated Capstone Concept [1] and are available online [2].

various threats and hazards. They are the most useful to commanders and decision makers and should be those that primarily inform their requirements. Lower-level metrics remain useful for comparing the performance and effectiveness of ISR capabilities, but they are limited to the *Sense* domain and do not necessarily take into account the various consequences that threats may have if authorities fail to react in a timely and appropriate fashion.

Value to Commander ↑

Category	Focus	Examples	
Measures of force effectiveness	Outcomes & risks (goes beyond ISR)	Depends on force's objectives · Risk of smuggling vessels reaching country	
Measures of ISR effectiveness	ISR effectiveness	Depends on ISR objectives · Probability of detecting vessels of type x in area y	
Picture quality metrics	Awareness & warning	· Completeness · Correctness · Timeliness	· Commonality · Extent / depth of information · Prediction perforfance
Fusion measures of performance	Data fusion	· Track accuracy · Track confidence · Track continuity	· False track ratio · Mean track life · Fusion latency
Sensor / sources measures of performance	Data collection	· Range / coverage · Persistence · Revisit / scan rate	· Accuracy · Reliability / error rate · Image quality

↑ Complexity of Assessment

Table 1: Examples of ISR metrics (adapted from [3])

1.1.5 Other analysis issues

There are many other issues pertaining to the analysis of ISR capabilities. An important one is the disconnect that sometimes occurs between force employers (e.g., operational commands) and force developers (e.g., acquisition and procurement organizations). Force developers tend to build capabilities for the future on the basis of planning scenarios, which may not completely or accurately reflect the threat spectrum faced by force employers. On the other hand, the analysis tools and methods used in support of capability development tend to be more sophisticated (e.g., synthetic simulation environments) than those available to operational commanders. That said, these tools are not necessarily better at assessing force effectiveness metrics and do not always leverage the expertise and lessons learned of the operators. Although force employers and force developers do not share the exact same goals and analysis requirements, the use of different analytical methods is rarely justified. The analysis performed in support of force employers should somehow inform part of the force development process.

1.2 Risk analysis and capability planning

Over the last decade, several risk assessment models have been developed and used by government organizations across the world to improve preparedness, response, or recovery from various threats, especially terrorist attacks. For instance, the U.S. Coast Guard's *Maritime Security Risk Analysis Model* (MSRAM) [4] has been deployed to every port in the U.S. and is now used to identify the main risks from potential terrorist attacks, and to allocate security resources accordingly. The *Transit Risk Assessment Methodology* (TRAM) [5] is used by the Department of Homeland Security (DHS) and applies a similar approach to identify risks to various public transportation systems. Models tailored to specific threat types (e.g., cyber attacks, biological attacks, chemical attacks) are also used by DHS and other government organizations to inform how their capabilities should be developed and allocated.

Most of these models are based on a conceptual framework that defines risk as a function of threat, vulnerability, and consequence (TVC). A recent review of various DHS risk assessment models by the National Research Council [6] found this construct generally appropriate for decomposing and organizing risk-related information. Although methodological issues have been identified with how some of these models implement it [6, 7], the TVC construct itself remains valid. It can be used to put into context the risks associated with different ISR vulnerabilities.

1.3 Aim

The aim of this paper is to demonstrate how risk analysis can inform decisions related to both the development and employment of ISR capabilities. It proposes a quantitative, all-hazards risk assessment model for comparing different ISR capability options or requirements, and prioritizing them. It uses risk as a primary metric and the basis of an objective function that can be minimized while setting operational and strategic capability requirements.

1.4 Outline

Following this introduction, Section 2 introduces the *Generalized Risk Assessment Model for Protection and Awareness* (GRAMPA), a Canadian model used here to illustrate how the overall risk-based approach described in this report can be implemented. Section 3 presents an example of how GRAMPA can be applied to improving maritime ISR capabilities in a domestic security context. Section 4 briefly describes how such model can be implemented and used to meet the complimentary needs of force employers and force developers in an integrated fashion. The main conclusions are summarized in Section 5.

In order to put the proposed approach into a practical context, some references are made to Canadian terminology and force structure throughout this paper. Nevertheless, the approach remains applicable to any defence or security organization. Examples of application to some scenarios of interest to Canadian authorities are presented, but in order to keep these examples unclassified, results are generated from *notional* data that do not accurately reflect real-world threats or capabilities.

2.0 RISK ASSESSMENT MODEL

2.1 Overview

GRAMPA is an all-hazards, quantitative risk assessment model. Like most of the models previously mentioned, it serves to compare the risks associated with different scenarios. These scenarios are broadly defined in the form of short vignettes that can reflect a variety of threats and hazards of interest to defence and security organizations. Table 2 shows a sample of an all-hazards risk assessment taxonomy [8] currently being developed by DRDC. It is used to classify the different types of scenarios of interest to federal partners in Canada, and to organize the scenarios analyzed in GRAMPA.

In GRAMPA, the TVC triad is decomposed into lower-level risk factors, including different vulnerability factors associated with ISR. This decomposition is done to a level that remains high enough to be broadly applicable to a wide variety of threats. GRAMPA itself does not evaluate individual threats, vulnerabilities, or consequences. These are assessed from other sources such as intelligence, historical analysis, operational performance data, modelling, or expert judgement. What GRAMPA provides is a means of centralizing, quantifying, aggregating, and visualizing the results of these assessments, while explicitly taking into account the uncertainties associated with each of them.

Category	Examples	
Man-made (intentional)	· Foreign state activity	· Organized criminal activity
	· Criminal activity (lone wolf)	· Terrorist activity
Man-made (unintentional)	· Spill	· Fire
	· Explosion	· Accident
	· Crash	· Collapse
Natural disaster	· Hurricane	· Storms
	· Landslide	· Flood
	· Fire	· Tsunami
	· Earthquake	· Volcano
Health disaster	· Pandemic	· Food contamination
	· Water contamination	

Table 2: All-hazards risk taxonomy (adapted from [8])

Uncertainties do not only exist in the assessments of the risk factors, but also in their interdependencies. These uncertainties should always be characterized, evaluated, and communicated in risk assessments. As mentioned in the National Research Council's review [6], a "proper recognition and characterization of both variability and uncertainty are important in all elements of a risk analysis, including effective interpretation of data as they are collected over time on threats, vulnerability, consequences, intelligence, and event occurrence".

In GRAMPA, the uncertainties associated with risk factors and their interdependencies are treated using fuzzy sets. Fuzzy set theory is increasingly used for risk analysis (e.g., [9–11]) as it is not only suited to deal with the aleatory uncertainty (variability) of the risk factors, but also the epistemic uncertainty (i.e., imprecision, vagueness, knowledge deficiency) surrounding them and their assessment through expert elicitation. The belief that an event will occur with probability p is described in terms of a membership function μ_p. In GRAMPA, every input has such membership function attached to it and it is characterized in the form of a triangular fuzzy number (TFN) defined by a minimum, most likely (ML), and maximum value, as shown in Figure 1. TFNs have the advantage of being simple and not too demanding to elicit from subject-matter experts, while still accounting for uncertainties in the inputs.

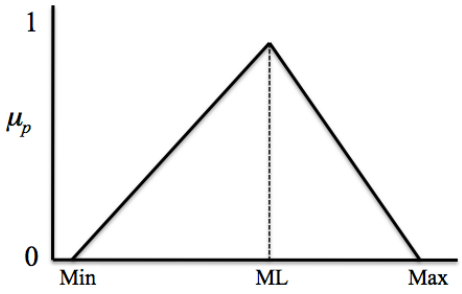

Figure 1: Triangular fuzzy number

GRAMPA is implemented in the form of a Microsoft Access® database with a user-friendly interface that facilitates the assessments of each TVC component and the communication of the results. The database provides a means of centralizing data and results for several scenarios or scenario variants.

2.2 Threat assessment

The *threat* is defined here as the expected frequency of occurrence of a particular type of potentially damaging event over the planning timeframe[2]. It is generally the most difficult risk factor to assess, especially for rare and intentional threats. For such threats, eliciting inputs from intelligence experts is often the only way of conducting the assessment. For frequent threats or hazards, quantitative analysis and modelling tend to be more appropriate than expert elicitation. Regardless of how the threat assessment is performed, a certain amount of uncertainty around it is normally expected.

[2]For force employers, the planning timeframe will generally be in the order of weeks or months. For force developers, the planning timeframe will generally be in the order of years.

In GRAMPA the threat is assessed on a six-point scale shown in Table 3 and Figure 2. Qualitative descriptions on the scale differ depending on the type of threat being assessed (intentional or non-intentional), following a practice of the intelligence community in Canada. A default frequency in the form of a TFN is associated with each threat level. However, the assessors *are not constrained to use these default frequencies*; any other TFN can be specified on the basis of modelling, expert judgement, or other information sources. This ability to deviate from default scales avoids potential losses in the resolution of the assessment, which occur when experts are forced to make a selection within pre-defined rating intervals or matrices [12, 13]. Another advantage of fuzzy assessments is that when different experts do not agree on a linguistic description or a particular frequency value, the TFN can be defined to span the full range of frequencies offered by the experts.

Description		Default frequency (expected events per year)		
(Intentional)	**(Non-intentional)**	**Min**	**ML**	**Max**
'Severe'	'Very Frequent'	10	55	100
'High'	'Frequent'	1	5.5	10
'Medium'	'Occasional'	0.1	0.55	1
'Low'	'Probable'	0.01	0.055	0.1
'Negligible'	'Improbable'	0.001	0.0055	0.01
'No Recognized Threat'	'Extremely Improbable'	0.0001	0.00055	0.001

Table 3: Threat assessment scales and default (modifiable) TFN values

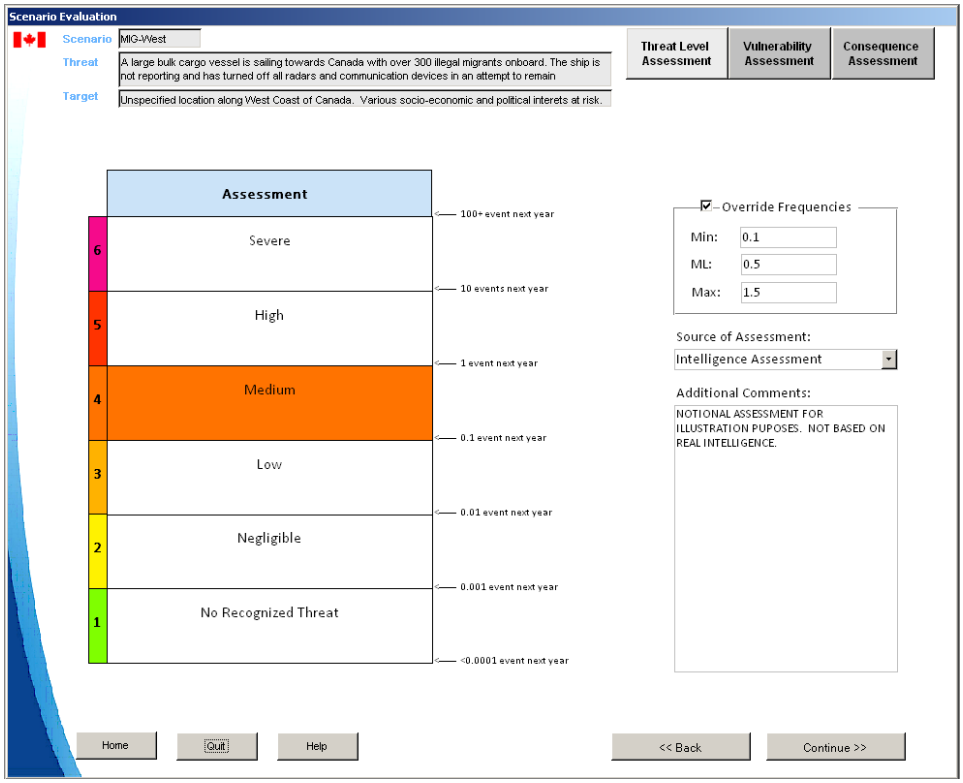

Figure 2: Threat assessment form in GRAMPA tool

Figure 2 provides a screen shot of the threat assessment form in GRAMPA. Default frequencies are assigned when the user selects a certain threat level on the scale, but these can be overridden by the user. In order to be applicable to both intentional and non-intentional threats, the assessment is not decomposed into intent and feasibility components. Unlike many other risk models, only the expected threat frequency is considered. However, this does not prevent the use of separate, specialized tools for estimating what the minimum, most likely, and maximum frequencies should be.

2.3 Vulnerability assessment

The *vulnerability* is defined here as the probability that the threat, assuming it exists and is capable of causing damage, will actually cause damage. It is a function of many factors, the most important being represented in the high-level fault tree of Figure 3. Essentially, damage can be prevented by deterring the threat, neutralizing it, or situating the target in such a way that the threat will not have harmful effects. The last of these can be achieved by either protecting the target, or making it inaccessible or unpredictable to the threat.

Neutralizing the threat requires many capabilities, including effective ISR capabilities. In essence, a threat must somehow be classified as such in a timely fashion, either through some anomaly detection (normally arising from surveillance and reconnaissance activities) or cueing (normally arising from intelligence activities)[3]. The threat also needs to be localized in a timely manner and its evolution (or trajectory) effectively tracked. Based on this, some kind of warning must be issued, and responsible authorities must be equipped, trained, and ready to neutralize or interdict the threat.

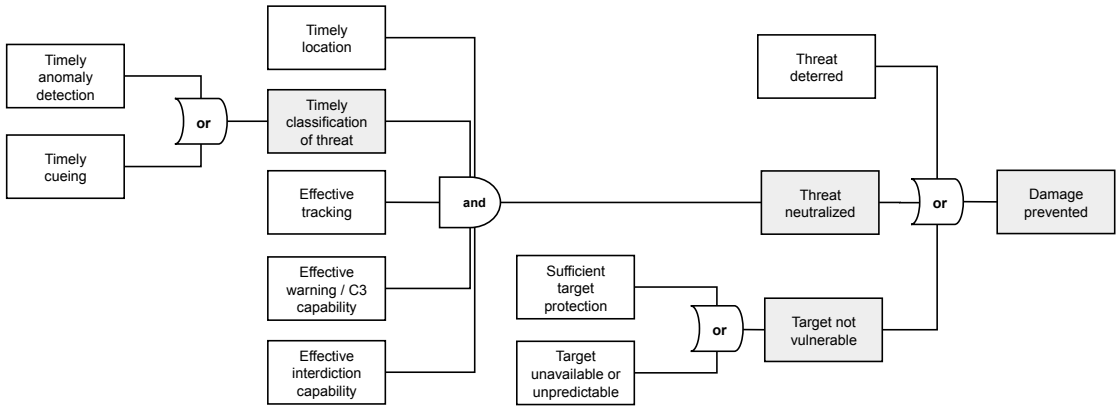

Figure 3: High-level fault tree for the vulnerability assessment

All these events do not have to occur sequentially. A threat may be flagged as such before it is localized, or its tracking may lead to a kinematic anomaly detection that flags it as a threat. The fault tree structure simply reflects, at a very high-level, the key requirements for preventing the threat from causing damage. This level of abstraction has the advantage that the fault tree is applicable to any type of threat or hazard, in any environment.

[3]The fault tree is deliberately framed differently than the traditional detection-classification-identification-tracking process, which is typically used to explain the process of generating and updating surveillance tracks within a common operating picture (COP). Here, detection refers to the ability to detect anomalies and classification refers to the ability to "flag" an entity as a threat, not merely obtaining its type or identity. Obviously, a good COP contributes to these tasks, but GRAMPA focuses on the higher-level outcomes of the ISR processes.

Some elements of the tree, such as the threat deterrence and target protection, are not decomposed further, but remain relevant to ISR requirements and should be considered. For instance, the act of conducting surveillance in a given region, or even simply demonstrating authorities' presence through patrols, may deter certain threats. Furthermore, some targets may be better protected or less accessible to certain threats depending on the situation or time of the year, which naturally reduces risks and should influence ISR requirements.

In GRAMPA, the probability of failure p_f associated with each input node of the tree (in white) is characterized by a TFN, in a similar way to the threat. Table 4 shows a qualitative scale that is used to assess each vulnerability component, with the corresponding default p_f values. Once again, this scale serves only to facilitate the assessment, but does not constrain it. Any p_f value between 0 and 1 can be specified, and TFNs may overlap with each other.

Description	Default p_f values		
	Min	ML	Max
'Certain to fail'	1	1	1
'Almost certain to fail'	0.85	0.95	0.99
'Probably will fail'	0.60	0.75	0.90
'About as likely as not to fail'	0.33	0.50	0.66
'Probably won't fail'	0.10	0.25	0.40
'Almost certain not to fail'	0.01	0.05	0.15
'Impossible to fail'	0	0	0

Table 4: Failure ratings and default (modifiable) TFN values

As for the threat assessment, input data can be obtained from different sources. Depending on the scenario analyzed, more advanced modelling can be performed outside of GRAMPA by further decomposing some elements of the fault tree. For instance, Boniface [14] breaks down the anomaly detection element into different types of anomalies or intelligence cues associated with maritime threats (e.g., vessel, cargo, passenger, crew, operator). Davenport [15], on the other hand, looks at kinematic anomalies and breaks them down by motion (speed, track, etc.), location (historical, current), and other factors. GRAMPA does not prescribe how the analysis should be done for each input node of the fault tree. For some scenarios, rough estimates obtained from expert elicitation or from interpreting recent experience or exercises may be sufficient.

The various components of the fault tree in Figure 3 may not be independent. For instance, in many scenarios, failing to localize the threat implies a failure in identification and tracking, since these capabilities may come from the same ISR assets. Accordingly, the level of dependence between capabilities needs to be considered. It is rarely possible to determine what this level of dependence is exactly, but as for the other inputs, a fuzzy assessment can be made.

Table 5 shows the rating scale, adapted from Ferdous [10], with the default coefficients of dependence that are currently implemented in GRAMPA in the forms of TFNs. For a given scenario, the level of dependence between fault tree components can be estimated using this scale. When necessary and possible, more accurate TFNs can be specified. Note that only positive dependence coefficients are considered here, but it would be easy to extend the model for considering negative dependence. Section 2.5 will explain how the coefficients are used in vulnerability calculations.

Level of dependence	Dependency coefficient D		
	Min	ML	Max
'Perfect independence'	0	0	0
'Very weak dependence'	0.005	0.10	0.20
'Weak dependence'	0.15	0.33	0.50
'Strong dependence'	0.45	0.66	0.85
'Very strong dependence'	0.80	0.90	0.995
'Perfect dependence'	1	1	1

Table 5: Interdependence scale and default TFN values

Figure 4 provides a snapshot of the vulnerability assessment screen in the GRAMPA tool. It shows how the different portions of the fault tree are coloured depending on the assessed p_f values. It also shows how the level of dependence can be specified for pairs of fault tree components.

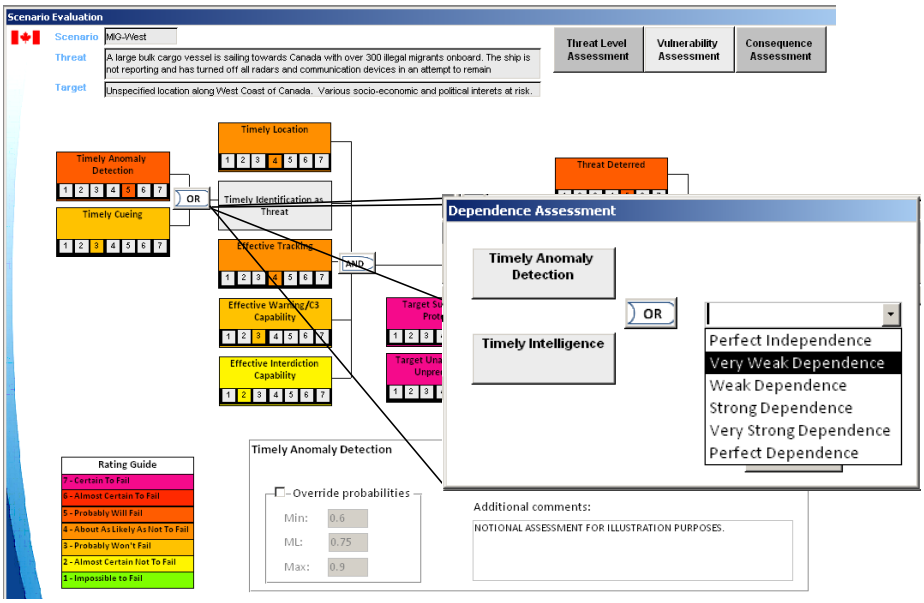

Figure 4: Vulnerability and dependence assessment forms in GRAMPA tool

2.4 Consequence assessment

The *consequence* assessment captures the various types of damage potentially inflicted by the threat and their magnitudes. Examples of consequence categories commonly used in all-hazards risk analysis and implemented in GRAMPA are shown in Table 6. Consequence categories that are relevant to stakeholders are identified at the onset of the risk analysis, and then assessed for each scenario. By default, each category is scored on a logarithmic scale shown in Table 7.

As with the previous risk factors, the assessment is not constrained by the rating scale. The default TFN values associated with each description can be modified, if necessary. The choice of consequence units does not matter, as long as it applies to all scenarios and consequence categories. If a single category of consequence is considered, then TFNs can be simply quantified in the units of the specific risks considered (e.g., number of fatalities, millions of dollars).

Consequence categories
Casualties
Direct economic losses
Secondary economic losses
Disruption of critical infrastructures / services
National / territorial security impact
Environmental impact
Political / reputation impact
Psycho-social impact

Table 6: Typical consequence categories

However, multiple consequence categories are normally considered in risk assessments. In order to derive an integrated risk score for ranking scenarios and requirements, the various consequences associated to a particular scenario must be summed up. This requires the use of common units for the different consequence scales. It also requires some kind of calibration between scales. For instance, if the units in Table 7 are interpreted as millions of dollars, some value judgement will be directly or indirectly required to apply this scale to other consequence categories, such as fatalities. Literature on the subject is reviewed in Jonkman [16]. Some risk assessment frameworks, such as those being developed for Public Safety Canada [8] or DHS [17], also provide scales that equate the magnitude of various

Description	Default values		
	Min	ML	Max
'Catastrophic'	1 000 000	5 500 000	10 000 000
'Extreme'	100 000	550 000	1 000 000
'Very severe'	10 000	55 000	100 000
'Severe'	1 000	5 500	10 000
'High'	100	550	1 000
'Moderate'	10	55	100
'Low'	1	5.5	10
'Very low'	0.1	0.55	1
'None or negligible'	0	0.055	0.1

Table 7: Consequence ratings and default (modifiable) TFN values

types of consequences. Any set of scales can be imported and used in GRAMPA, as long as the scales and their calibration are acceptable to all stakeholders.

2.5 Risk Assessment

2.5.1 Manipulating TFNs

The previous sections show how various TVC components can be estimated and assessed in the form of TFNs. These are easy to define and understand, while capturing the uncertainties. Arithmetic operations on them are also fast and easy to compute. The result of adding or subtracting two TFNs, for instance $\tilde{A} = (a_{\min}, a_M, a_{\max})$ and $\tilde{B} = (b_{\min}, b_M, b_{\max})$, is also a TFN:

$$\tilde{A} + \tilde{B} = (a_{\min} + b_{\min}, a_M + b_M, a_{\max} + b_{\max}) \quad (1)$$

$$\tilde{A} - \tilde{B} = (a_{\min} - b_{\min}, a_M - b_M, a_{\max} - b_{\max}) \quad (2)$$

The product of two TFNs is not a TFN, but it is generally approximated as such in order to simplify calculations [9]:

$$\tilde{A} \times \tilde{B} \approx (a_{\min} b_{\min}, a_M b_M, a_{\max} b_{\max}) \quad (3)$$

Furthermore, it is common to reduce the interval covered by TFNs using α-cuts. Simply put, α-cuts convert fuzzy intervals into crisp (or "defuzzified") intervals by specifying a minimum threshold (α) for the membership function. In other words, they can be used to establish a minimum "degree of confidence" for the parameter of interest within that interval. Figure 5 shows how this is done for a particular TFN where $p_{\min}$, p_M, and $p_{\max}$ represent the minimum, most likely, and maximum values, respectively, and μ_p is the degree of membership (ranging between 0 and 1) of the different probabilities.

The triangular fuzzy number $\tilde{P}^{\alpha}$ resulting from the α-cut is obtained through Equation 4. It spans the interval between p_L^{α} and p_R^{α}, and corresponds to the shaded area in Figure 5.

$$\begin{aligned} \tilde{P}^{\alpha} &= (p_L^{\alpha}, p_M, p_R^{\alpha}) \\ &= (p_{\min} + \alpha(p_M - p_{\min}), p_M, p_{\max} - \alpha(p_{\max} - p_M)) \end{aligned} \quad (4)$$

Now assume that a TFN has already been defined for all inputs to a fault tree such as the one in Figure 3. Based on Equation 3, the output of a logic gate can be approximated as a TFN. The operations of Table 8 serve to determine the α-cut TFN (p_L^{α}, p_M,p_R^{α}) resulting from an "OR" or "AND" gate with n inputs.

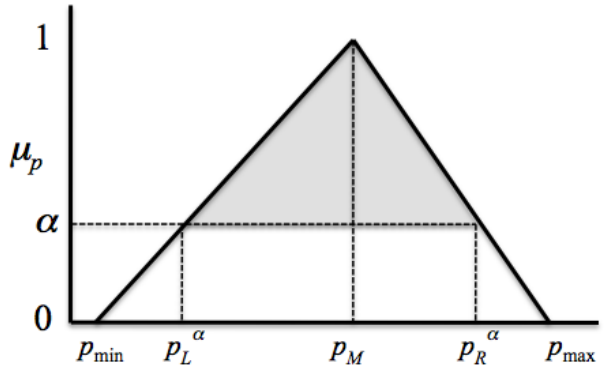

Figure 5: Triangular fuzzy number and α-cut interval

Operation	Formulation		
"OR" gate	$p_L^{\alpha} = 1 - \prod_{i=1}^{n}(1 - p_{iL}^{\alpha})$;	$p_M = 1 - \prod_{i=1}^{n}(1 - p_{iM})$;	$p_R^{\alpha} = 1 - \prod_{i=1}^{n}(1 - p_{iR}^{\alpha})$
"AND" gate	$p_L^{\alpha} = \prod_{i=1}^{n} p_{iL}^{\alpha}$;	$p_M = \prod_{i=1}^{n} p_{iM}$;	$p_R^{\alpha} = \prod_{i=1}^{n} p_{iR}^{\alpha}$

Table 8: Traditional α-cut-based fuzzy logic operations [9, 10]

These operations assume that inputs are perfectly independent from each other, which may not be the case. Given two inputs $\tilde{P}_a$ and $\tilde{P}_b$, and a fuzzy dependence coefficient $\tilde{D}$ between them, as defined in Table 5, a different formulation for the "AND" and "OR" logic operations derived from the Frank family of copula [18] is used. Table 9 shows equations for deriving p_L^{α}, where $s = \tan(\pi(1 - D_L)/4)$. Similar operations can be used for deriving p_M and p_R^{α}.

Operation	Formulation	
"OR" gate	$p_L^{\alpha} = \begin{cases} 1 - (1 - p_a^{\alpha})(1 - p_b^{\alpha}) & \text{if } D_L = 0 \\ \max(p_a^{\alpha}, p_b^{\alpha}) & \text{if } D_L = 1 \\ 1 - \log_s \left[1 + (s^{1-p_a^{\alpha}} - 1)(s^{1-p_b^{\alpha}} - 1)/(s-1)\right] & \text{if } 0 < D_L < 1 \end{cases}$	
"AND" gate	$p_L^{\alpha} = \begin{cases} p_{aL}^{\alpha} p_{bL}^{\alpha} & \text{if } D_L = 0 \\ \min(p_{aL}^{\alpha}, p_{bL}^{\alpha}) & \text{if } D_L = 1 \\ \log_s \left[1 + (s^{p_{aL}^{\alpha}} - 1)(s^{p_{bL}^{\alpha}} - 1)/(s-1)\right] & \text{if } 0 < D_L < 1 \end{cases}$	

Table 9: Modified α-cut-based fuzzy logic operations with dependence coefficient D [19]

This formulation and other possible models for combining dependent variables are reviewed at Ref. [19]. The functions presented in Table 9 are continuous. Note that for $D = 0$, the same results as in Table 8 are obtained. For logic gates with more than two inputs, the "AND" and "OR" can be performed sequentially, starting by the first pair of inputs, combining the result to the third input, and so on.

2.5.2 Overall risk assessment

Given the fuzzy assessments made for each input to the fault tree, the operations above are performed to derive the overall probability of failure, or vulnerability $\tilde{V}_\alpha$, associated to a particular scenario. Since the threat assessment here is already a single TFN, the overall frequency of failure with which a threat is expected to cause damage is the product $\tilde{T}_\alpha \times \tilde{V}_\alpha$ (calculated as per Eq. 3) of the α-cut threat and vulnerability. The overall consequence $\tilde{C}_\alpha$ is obtained by summing up the α-cut TFNs associated with n individual consequence categories, as shown in Equation 5.

$$\tilde{C}_\alpha = \Big(\sum_{i=1}^{n} c_{iL}^{\alpha}, \sum_{i=1}^{n} c_{iM}, \sum_{i=1}^{n} c_{iR}^{\alpha} \Big) \tag{5}$$

The overall risk $\tilde{R}_\alpha$ is typically calculated as $\tilde{R}_\alpha = \tilde{T}_\alpha \times \tilde{V}_\alpha \times \tilde{C}_\alpha$, but other risk functions are possible[4]. The straight product has the advantage of being simple to explain and compute, but it does not convey any aversion of the decision makers to highly consequential events; a threat expected to cause damage with a certain frequency will have the same risk score as a threat expected to cause damage half as frequently, but with twice the consequences each time[5]. Implicitly here, the disutility is assumed to increase linearly with the consequences. Various ways of integrating risk aversion into the risk function can be found at Ref. [16].

2.5.3 Adversarial intent

Currently, GRAMPA does not account for potential dependencies between TVC components in the risk function. Such dependencies may exist for intentional threats, especially terrorists, who are arguably more likely to strike where they perceive higher vulnerabilities and consequences. For these threats, it is assumed that the threat assessment elicited from subject-matter experts will implicitly scale the adversarial intent considering perceived vulnerabilities and consequences. One way to make this more explicit in the model would be to break intentional threat assessments into two components: intent and capability. These could be assessed separately in the form of TFNs, but the assessments would still require reliable estimates of potential attackers' intents, perceptions, capabilities, and decisions, which are all subject to large uncertainties. A number of techniques and models have recently been developed for counterterrorism decision making [20] and could inform such assessments as GRAMPA is further developed.

Furthermore, the fact that dependencies between TVC components are not explicitly modelled is not an issue when these components are presented separately to decision makers. As will be shown in the following example, TVC components do not necessarily have to be integrated into a single risk score in order to compare different scenarios. Although a single score may be useful for ranking different scenarios or performing certain optimization tasks, decision makers are (and should be) presented with more information than a single risk score when comparing scenarios and capability options.

[4]Since risk values can span several orders of magnitude, it may be convenient to represent them in those terms. In GRAMPA, the risk can be scored on a logarithmic scale ranging from zero to nine, using the following equation:

$$\tilde{R}_\alpha = 9.0 \times \frac{\log\Big[(\tilde{T}_\alpha \times \tilde{V}_\alpha \times \tilde{C}_\alpha + \epsilon)/\epsilon\Big]}{\log\Big[(T_{\max} \times V_{\max} \times C_{\max} + \epsilon)/\epsilon\Big]} \tag{6}$$

where $T_{\max}$, $V_{\max}$, and $C_{\max}$ are constants representing the highest possible values that T, V, C can take in the model. ϵ is a very small positive number used to keep the score positive when $\tilde{T}_\alpha \times \tilde{V}_\alpha \times \tilde{C}_\alpha$ is smaller than one or equal to zero.

[5]Some "risk-averse" decision makers would see more disutility in the second scenario and would be more inclined to prevent it.

3.0 EXAMPLE

This section is divided into three parts and illustrates how the approach can be used. The first part demonstrates how a specific scenario is assessed using GRAMPA. The second part establishes a notional "baseline" risk profile for a few selected maritime security scenarios (illegal migration, illegal fishing, and pollution) at different locations along Canadian coasts (West Coast, East Coast, Arctic). The third part compares the amount of risk mitigation provided by introducing different capability options. Note that in order to keep this paper unclassified, all the results presented here have been generated from **fictitious input data** that deliberately do **not** accurately represent TVC components.

3.1 Risk analysis - single scenario

Consider the following scenario:

> *"A large bulk cargo vessel is sailing towards Canada with over 300 illegal migrants onboard. The ship is not reporting and has turned off all radars and communication devices in an attempt to remain undetected."*

A first variant of this scenario, hereby referred to as 'MIG-West', will assume that the ship departed from Asia and is heading toward an unknown destination along the west coast of Canada. For the purpose of this example, it is assumed that this threat has been assessed as 'Medium' from intelligence sources, with an expected frequency of $\tilde{T} = (0.1, 0.5, 1.5)$ vessels per year over the planning timeframe considered by the analysis.

For the vulnerability assessment, the timeliness of the events must be well understood, since a warning must be issued early enough to enable an adequate response by responsible authorities. To analyze this timeliness issue, analysts from the North American Aerospace Defense Command (NORAD) have developed a deterministic model called the Maritime-Timeline Analysis and Requirements Toolset (M-TART) [21].

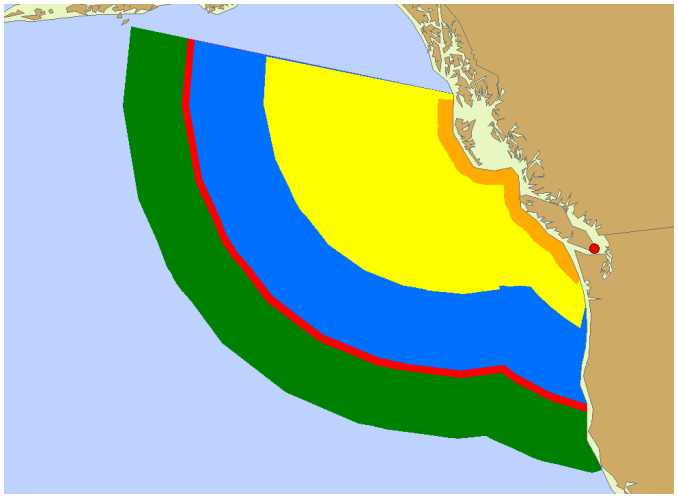

Figure 6: M-TART results for MIG-West scenario (from notional response capability data)

In the present example, the destination of the migrant ship is uncertain and it could be bound for anywhere along the Canadian west coast. Given the location of response assets (Victoria, BC in this example, shown by the red dot in Figure 6), M-TART determines the boundary before which a warning must be issued. As shown in Figure 6, in order to intercept the vessel at the start of the orange zone, the response ship must start sailing when the threat crosses into the yellow zone. That implies that a sailing order must have been issued before the migrant ship enters the blue zone, and that the first warning of the threat must have been issued by the time the migrant vessel enters the red zone. Obviously, the size of these zones would vary depending on the speed of the threat, the speed of the response assets and their notice to move[6].

[6] Although M-TART has been developed to help military planners determine maritime warning requirements, similar timeline analyses have been generated for the air domain, and other domains could be analyzed in the same way.

The next step consists of estimating the probability that an anomaly or cueing would trigger a warning before the migrant ship reaches the red zone. Once again, there are tools available to assist with obtaining a realistic estimate. For example, the Surveillance Analysis Workbook (SAW) [22] has been developed for helping CF maritime commanders plan and assess surveillance activities along Canada's coasts. It takes into account the threat size and speed, as well as revisit times of different surveillance assets (sea-, land-, air-, space-based) and their individual performance in terms of detection [23, 24]. The tool is currently used to produce weekly and quarterly statistics for surveillance effectiveness along the Canadian coasts. These statistics can be used to estimate probabilities of detecting and identifying threats from surveillance activities.

In the present example, other risk components are assessed by various means, including historical records and expert judgment. In some cases, detailed modelling is either not required to obtain a sufficient level of precision, or simply not possible with available data and resources. Notional assessments of the risk factors for the MIG-West scenario are presented in Table 10.

Factor	Rating	Min	ML	Max	Comments
Threat	'Medium'	0.1	0.5	1.5	Notional assessment from intelligence and trend analysis.
Timely anomaly detection	'Probably will fail'	0.60	0.75	0.90	Notional assessment from M-TART. Results indicate that warning must occur before the vessel reaches the red zone of Fig 6. Anomalous non-reporting will be noticed before that point only by aerial surveillance. Notional results from SAW indicate that such surveillance is infrequent in this area and will likely fail at detecting anomaly.
Timely cueing	'Probably won't fail'	0.10	0.25	0.40	Notional assessment from intelligence. Efforts devoted to migrant networks would likely provide cueing before the vessel reaches the red zone of Fig 6, but some smuggling ventures may not be known.
Timely localization	'About as likely as not to fail'	0.33	0.50	0.66	Notional assessment from modelling. Target should be localized before response assets start to sail, i.e., before the threat reaches the yellow zone of Fig 6. Notional results from SAW help to estimate probability of failure.
Effective tracking	'Probably won't fail'	0.10	0.25	0.40	Notional estimates from modelling. Air and space-based assets can provide adequate track update rates in this region.
Effective C3 and warning	'Almost certain not to fail'	0.01	0.05	0.15	Notional estimates based on historical experience. Maritime command, control, communications (C3) and warning procedures are in place. Response plans are routinely exercised and have been successfully used.
Effective interdiction	'Almost certain not to fail'	0.01	0.05	0.15	Notional estimates based on historical experience. Ready duty ship nearly always available and backup ships available when needed. Migrant ships successfully interdicted in recent years.
Deterrence	'Probably will fail'	0.60	0.75	0.90	Notional assessment from historical experience. Threat level has not significantly changed despite the fact that authorities have thwarted many smuggling operations.
Target protection	'Almost certain to fail'	0.85	0.95	0.99	Notional assessment from expert judgement. The 'target' here is not specific. Ship reaching any point of the coast would represent a failure. Small chance that response assets will already be on target's trajectory by coincidence.
Target availability	'Almost certain to fail'	0.85	0.95	0.99	Notional assessment from expert judgement. The 'target' in this case cannot be removed or made unpredictable. Small chance that sea conditions will prevent ocean crossing.
Casualties	'Low'	1	5.5	10	Notional assessment from expert judgement. Conditions onboard ship and risks associated with boarding may lead to casualties.
Economic losses	'Moderate'	10	55	100	Notional assessment from historical experience. Disposal cost of a recent migrant ship on the order of $25M. Cost associated with processing of migrants and prosecution of organizers of the same magnitude.
Infrastructure disruption	'None or negligible'	0.01	0.05	0.1	Notional assessment from expert judgement. Negligible impact on infrastructure expected.
Environmental damage	'Very low'	0.1	0.55	1	Notional assessment from expert judgement. Minor environmental damage possible due to poor ship conditions but would be easily remediated.
National security	'Low'	1	5.5	10	Notional assessment from expert judgement. Ship may carry suspected terrorists that will require investigation by law enforcement agencies.
Reputation and influence	'Moderate'	10	55	100	Notional assessment from expert judgement. Absence of an appropriate response could damage Canada's reputation and some international relations.

Table 10: Notional assessment of scenario MIG-West for illustrative purposes

If, for simplicity, it is assumed that the dependence coefficient is $D = 0$ everywhere, and if an α-cut of 0.05 is applied to TFNs, the resulting values are $\tilde{T}_\alpha = (0.12, 0.50, 1.45)$, $\tilde{V}_\alpha = (0.24, 0.49, 0.75)$, and $\tilde{C}_\alpha = (27.1, 121.6, 216.1)$. Their overall product is $\tilde{R}_\alpha = (0.77, 29.8, 234)$, which means that the risk associated with the MIG-West scenario is believed to be equivalent to a figure somewhere between \$0.8M and \$234M, with a most-likely value of \$30M. This range is quite wide, spanning two orders of magnitude, but it does reflect the wide uncertainties associated with some of the risk components. A narrower range could be obtained by conducting further analysis on some of the parameters. Before doing so, however, it is preferable to see how this range compares to the risk associated with other scenarios of interest.

3.2 Baseline risk profile for multiple scenarios

Figure 7 presents a frequency-consequence plot that includes the MIG-West scenario and analogues for the east (MIG-East) and Arctic (MIG-Arct) coasts. It also contains notional results for all three coasts for two other maritime scenarios: one where a ship is fishing illegally (FSH-x) and one where a ship is illegally discharging oil while en route (POL-x). Variants of these scenarios for all three coasts of Canada are presented. The risk values are shown with error bars reflecting the α-cut intervals for the frequency and consequence assessments. The diagonal lines represent iso-risk values.

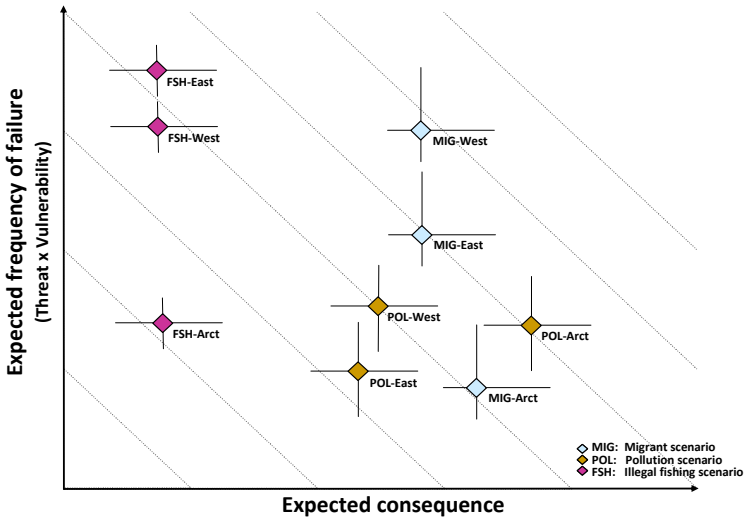

Figure 7: Notional baseline risk profile for selected scenarios

This plot shows how, by using risk as a metric, multiple scenarios of very different natures can be directly compared. Other scenarios not necessarily related to the maritime domain could also be added.

3.3 Options analysis and risk management

Once a baseline risk profile has been established, various options related to the employment and development of capabilities can be explored. The beneficial (or detrimental) impact of some capability options on the baseline risk profile are shown in Figure 8. For example:

- Re-allocating existing aerial surveillance hours to the West Coast (Fig. 8a) may mitigate risks for scenarios in that region, e.g., MIG-West, but may increase the risk associated with threats in other regions.

- Increasing intelligence effort devoted to migrant networks (Fig. 8b) may reduce the risk associated with that type of threat, but may not have any effect on the risk associated with other threats.
- The acquisition of a new space-based surveillance capability (Fig. 8c) may have a beneficial impact on all scenarios, albeit to different extents.
- Improving the response force posture (Fig. 8d), either by changing its location, notice to move, or the type of response capability involved, may also have a generally positive impact. As can be demonstrated using M-TART, improving the response force posture can reduce ISR requirements over time and space. Accordingly, even if no additional ISR assets are introduced, an improved response posture may make ISR requirements easier to meet with available resources. At the same time, an improved response posture or capability may reduce the potential consequences associated with various scenarios.

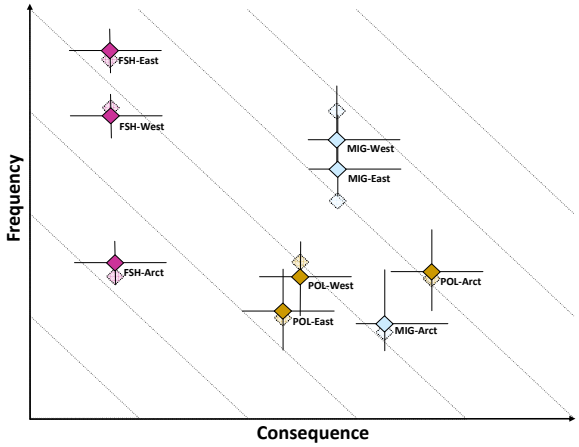

(a) Aerial patrol hours re-allocated to West Coast

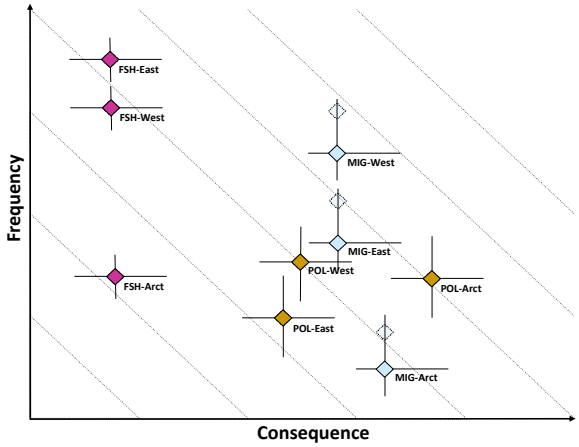

(b) Increased intelligence effort on migrant networks

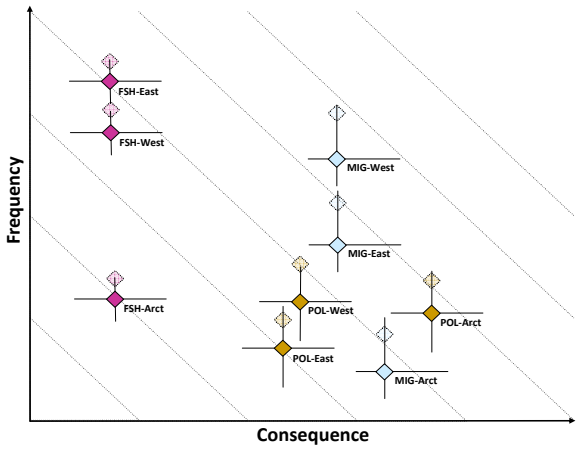

(c) Additional space-based surveillance capability

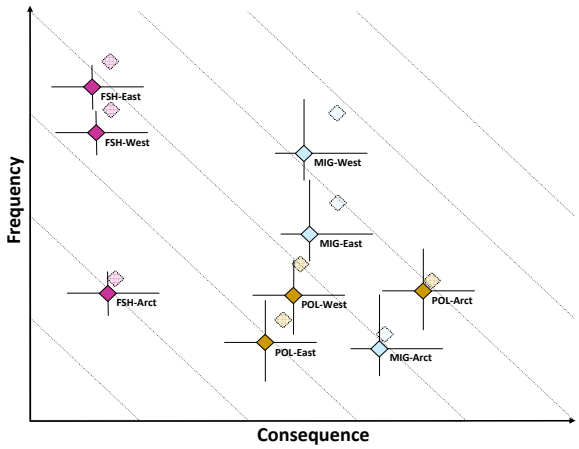

(d) Improved response posture

Figure 8: Notional risk profile modified for different capability options. Shaded points correspond to the pre-mitigation baseline of Figure 7.

This example demonstrates the importance of understanding the *global* risk-mitigation effect of different capability options. As identified by Cox [25], global risk mitigation is not simply a matter of allocating resources to scenarios presenting the highest risk (MIG-West in our example). The correlation between different scenarios must be understood and exploited in identifying optimal risk-reducing options. In other words, it is the overall risk-mitigation impact that should be considered in determining where force employment and force development resources should be invested.

4.0 IMPLEMENTATION

Perhaps the most challenging aspect of such an approach to capability analysis lies in how it is implemented and applied by organizations. As mentioned in the introduction, there are generally many stakeholders involved in the employment and development of ISR capabilities, even within a single organization. Who should own or contribute to the planning and analysis processes may be a contentious issue. Nevertheless, force employers and force developers can take advantage of a risk-based approach, and this approach can be applied in an integrated fashion.

Force employers will primarily set their requirements based on operational imperatives. By leveraging threat and vulnerability assessments already conducted for operations using a tool such as GRAMPA, they can find ways of augmenting the effectiveness of resources already available to them, by allocating them in a way that minimizes the risks associated with their mission set. As a secondary function, force employers can also use this approach to identify, in a quantifiable and defensible manner, the risks posed by current and future threats that they cannot mitigate to a level that they deem acceptable. In turn, these results can be used to justify new capability requirements and to advocate for additional capabilities or resources.

Force developers can then use the same risk-based approach for comparing different capability options and identifying which ones best meet the requirements set by force employers, as well as the long-term requirements of the organization. Since all governments operate in fiscally constrained environments, using a risk-based assessment process to inform resource allocation and acquisition allows departments to use public funds in a responsible and defensible manner.

Having all stakeholders using the same risk-based approach has other advantages. Even though there may be differing opinions on the scenarios to be used and the assessments of the individual risk components, using a common framework can facilitate discussions on capability requirements. Inevitably, some differences of opinion will arise during the risk assessment, especially at the time of quantifying some of the risk factors. GRAMPA alleviates this problem, since it does not require a single value for each risk factor, but rather a fuzzy interval that can encompass estimates from different experts or organizations. Another way to alleviate conflicting opinions is to favour the use of inputs backed up by sound quantitative analysis. Using operationally accepted modelling tools such as M-TART and SAW is an example of this.

Additionally, because the approach is scalable, it can be used on a regional basis to assess the risks associated with specific areas of responsibility, and to inform decisions associated with tactical and operational planning. Assuming that the same framework is employed in every region, the results can be compiled and compared at a national level, or even on a pan-departmental basis, in order to inform strategic planning and policies. This kind of multi-tiered approach has been successfully used by the U.S. Coast Guard for the application of MSRAM to port security issues [4].

5.0 CONCLUSION

The approach presented here provides a systematic and quantitative way of comparing the risk mitigation effects resulting from different ISR capability options. At an operational level, these options could be new ways of employing or allocating available resources. At a strategic level, the approach could be used to inform not only resource allocation decisions, but also decisions related to policy and the procurement of new ISR systems. Hence the approach could inform both the development and employment of ISR capabilities.

The main advantage of the approach is that it can be applied relatively quickly to a wide range of scenarios and scenario variants. Risk factors can be initially estimated through expert elicitation and, where necessary, refined through quantitative analysis and modelling in order to reduce their uncertainties. Regardless of how the inputs are derived, the uncertainties are explicitly considered in GRAMPA – an important feature that traditional matrix-based or interval-based risk models generally lack.

Another benefit of the approach presented here is that the analysis is not limited to *Sense* functions. It does take into account, to a certain extent, functions related to the *Command*, *Shield*, and *Act* domains. As such, it provides a level of integration that does not always exist in other ISR capability planning approaches. It also bridges the gap between purely threat-based and purely capability-based (or vulnerability-based) approaches to ISR analysis. Although the examples presented here focus on the maritime domain, scenarios pertaining to other environments could be analyzed in the same way and compared using the same risk measures, allowing for a rigorous assessment of ISR options across all domains. Additionally, this approach could be used at the regional and national levels, and possibly on an interdepartmental basis, in an effort to generate a unified, cohesive risk-assessment product, and to inform decision makers at all levels.

ACKNOWLEDGMENTS

The authors would like to thank Dr. Patrick Dooley for his valuable comments and suggestions during the course of this work.

REFERENCES

[1] Department of National Defence, “Integrated Capstone Concept,” ISBN 978-1-100-16456-4, 2009.

[2] Department of National Defence, “Capability Domains - Definitions,” 2011, http://www.cfd-cdf.forces.gc.ca/sites/page-eng.asp?page=6433 (Last accessed 09 Oct 2011).

[3] Gauthier, Y., Bourdon, S., Dore, S., and Fong, V., “Defining and selecting metrics for intelligence, surveillance and reconnaissance,” DRDC CORA DOR(MLA) Research Note RN 2004/08, 2004.

[4] Fu, J., “Maritime Security Risk Analysis Model,” Proceedings of USCG-CREATE Maritime Risk Symposium, Nov. 2010, http://create.usc.edu/Fu-Mowrer%20-%20Pres.pdf (Last accessed 20 Aug 2011).

[5] Security Analysis and Risk Management Association, “Transit Risk Assessment Methodology (TRAM),” http://www.sarma-wiki.org/index.php?title=Transit_Risk_Assessment_Methodology_(TRAM) (Last accessed 23 Nov 2011).

[6] National Research Council, “Review of the Department of Homeland Security’s Approach to Risk Analysis,” National Academies Press, 2010, http://www.nap.edu/catalog.php?record_id=12972 (Last accessed 20 Aug 2011).

[7] Szwed, P., Boniface, D., and Myers, J., “An examination of risk-based resource allocation in the national strategy for homeland security using the maritime domain as context,” Proceedings of the American Nuclear Society - 2005 International Topical Meeting on Probabilistic Safety Analysis, 2005.

[8] Goudreau, A. and Verga, S., "All-Hazards Risk Assessment Methodology," Update for DRDC DSTIC Risk Analysis Workshop, 2011.

[9] Tyagi, S., Pandey, D., and Tyagi, R., "Fuzzy set theoretic approach to fault tree analysis," International Journal of Engineering, Science and Technology, Vol. 2, No. 5, pp. 276-283, 2010, http://www.ajol.info/index.php/ijest/article/viewFile/60165/48415 (Last accessed 20 Aug 2011).

[10] Ferdous, R., Khan, F., Sadiq, R., Amyotte, P., and Veitch, B., "Fault and event tree analyses for process systems risk analysis: uncertainty handling formulations," Risk Analysis, Vol. 31, No. 1, 2011.

[11] Bourdon, S., Fong, V., and Caron, J.-D., "A Capability-Based Risk Assessment Methodology - Application to the Prioritization of Capability Gaps at NORAD and USNORTHCOM," DRDC CORA Technical Report TR2010-172, 2010.

[12] Hubbard, D. and Evans, D., "Problems with scoring methods and ordinal scales in risk assessment," IBM J. Res. & Dev., Vol. 54, No. 2, Jun 2010.

[13] Cox, L., "What's Wrong with Risk Matrices?" Risk Analysis, Vol. 28, No. 2, 2008, http://risksociety.org.nz/file-uploads/risk%20matrix.pdf (Last accessed 20 Aug 2011).

[14] Boniface, D., "Risk-Based Resource Allocation in Maritime Security and Maritime Domain Awareness," Proceedings of MORS Meeting on Analytical Support to Maritime Domain Awareness and Counter Piracy, Oct 2009, http://www.mors.org/events/mda_presentations.aspx (Last accessed 20 Aug 2011).

[15] Davenport, M. and Carson, N., "Kinematic Behaviour Anomaly Detection - Final Report," DRDC CORA Contractor Report CR2008-002, Apr 2008.

[16] Jonkman, S., van Gelder, P., and Vrijling, J., "An overview of quantitative risk measures for loss of life and economic damage," Journal of Hazardous Materials A99 pp. 1-30, 2003.

[17] Keeney, R. and Winterfeldt, D., "A Value Model for Evaluating Homeland Security Decisions," Risk Analysis, Vol. 31, No. 9, 2011.

[18] Frank, M., "On the simultaneous associativity of and F(x, y) and x+y-F(x, y)," Aequationes Mathematicae Vol. 19 pp. 194-226, 2009.

[19] Ferson, S. et al, "Dependence in probabilistic modeling, Dempster-Shafer theory and probability bounds analysis," Technical Report SAND 2004-3072, Sandia National Laboratories, 2004, http://www.ramas.com/depend.pdf (Last accessed 13 Nov 2011).

[20] Merrick, J. and Parnell, G. S., "A comparative analysis of PRA and intelligent adversary methods for counterterrorism risk management," Risk Analysis, Vol. 31, No. 9, 2011.

[21] Carson, N. and Caron, J.-D., "The Maritime Timeline Analysis and Requirements Toolset (M-TART)," PHALANX, Dec 2010, http://dodreports.com/pdf/ada538828.pdf (Last accessed 20 Aug 2011).

[22] Wind, A. and Horn, S., "Surveillance Analysis Workbook (SAW) v11.2," DRDC CORA Technical Memorandum TM2009-039, Aug 2009.

[23] Horn, S., Carson, N., and Wind, A., "A Metric for Maritime Intelligence, Surveillance, and Reconnaissance (ISR) Probability of Identification," DRDC CORA Technical Memorandum TM2009-037, Sep 2009.

[24] Horn, S., "Recognized Maritime Picture Tools and Analysis," Proceedings of MORS Workshop on Maritime Domain Awareness, Oct 2009, http://www.mors.org/events/mda_presentations.aspx (Last accessed 20 Aug 2011).

[25] Cox, L., "What's Wrong with Hazard-Ranking Systems: An Expository Note," Risk Analysis, Vol. 29, No. 7, 2009.