Provenance Tracking of Hybrid Decision Making

Douglas S. Lange **Crisrael Lucero** **Braulio Coronado**
Naval Information Warfare Center, Pacific
UNITED STATES OF AMERICA

dlange@niwc.navy.mil

ABSTRACT

Recording the provenance of all information changes within a system or integrated systems of systems provides vital information about the decisions being made and the situations that prompted those decisions. From a forensics standpoint, this can be used to recreate the decision environment. However, provenance can also serve two other important functions. The data collected can support integration of components, and the graph data structure generated can support operator situational awareness through explanation, summarization, and alerting.

Hybrid warfare will necessarily bring together disparate decision support capabilities as the decision makers must operate in multiple warfare domains. Autonomous agents will likely play roles in the planning and execution processes, at times being enabled to make decisions without human intervention, but that human decision makers must become aware of. Provenance graphs have been shown to be translatable to Rhetorical Structure Graphs (RSG) allowing agents to explain their actions to humans in natural language and even with multimodal communication. Provenance has also been shown to enhance the monitoring of plan execution, and can be used to provide notification to human or autonomous agent that a plan might need to be rethought when information that was used in the planning has changed. As we move towards intelligent machines working to support teams of human decision makers in complex environments, the need to track decisions and their input becomes vital.

INTRODUCTION TO PROVENANCE

Provenance is information about entities, activities, agents, and the relationships between these concepts [1]. This information explains more than what happened, it also answers questions about how the entities were manipulated, when that occurred, and who was involved in the process. We are likely familiar with news and fictional stories about tracking the provenance of art works. The provenance of the creation, destruction, or modification of any entity can be tracked. In this paper, we will focus on the information held within military systems. Within Command and Control (C2), information provenance is necessary for recording the decision-making process behind actions, especially when autonomous and artificial intelligence (AI) agents are deeply involved. The “who” that might be involved in a process may be a human or an AI agent.

Information provenance serves several purposes. Forensically, provenance traces provide the people and agents involved in making decisions and how data has evolved to get to that decision. The Association of Computing Machinery U.S. Public Policy Council stated that data provenance is an explicit principle for algorithmic transparency and accountability [2]. Fully recorded provenance can shed light on data dependencies, responsibility flow, and help explain why certain actions were taken. As AI and autonomous agents continue to automate processes, they have become more integral in making critical decisions [3].

The PROV-DM Model

The Provenance Data Model (PROV-DM) was developed by the World Wide Web Consortium (W3C) and defines a universal representation of information provenance [1]. *Entities* are things that can be interacted with through *activities*, which are performed by *agents*. Relationships between the PROV-DM descriptions are used to model the dependencies and transformations of data as it changes over time. Although the building blocks in PROV-DM are relatively simple and generic, they can be used to express highly complex system representations. The PROV-DM is computationally represented as directed acyclic graphs where *entities*, *activities*, and *agents* are *nodes* and the relationships between those three core concepts are *edges*.

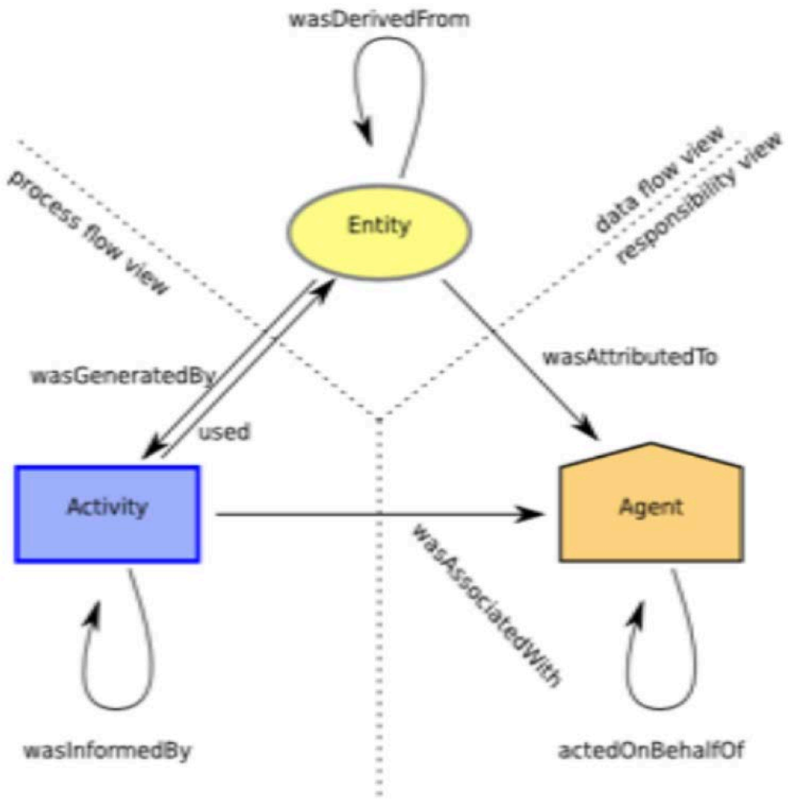

Figure 1 The PROV-DM Model. [1] Used with permission by the author.

Decision Support Based on Provenance

Using this simple model, several capabilities have been demonstrated using a prototype command and control system. The Intelligent Multi-uxv Planner with Adaptive Collaborative/control Technologies (IMPACT) was developed to experiment with the idea of allowing humans and AI agents to combine in order to allow multiple teams of systems with varying degrees of autonomy to be directed from a single workstation. Rather than many people being used to operate a single unmanned vehicle, several heterogeneous vehicle teams could be managed by one person [3]. In one set of experiments, IMPACT was enhanced by capabilities developed by allied countries: some capabilities were integrated within, and some interfaced across system boundaries.

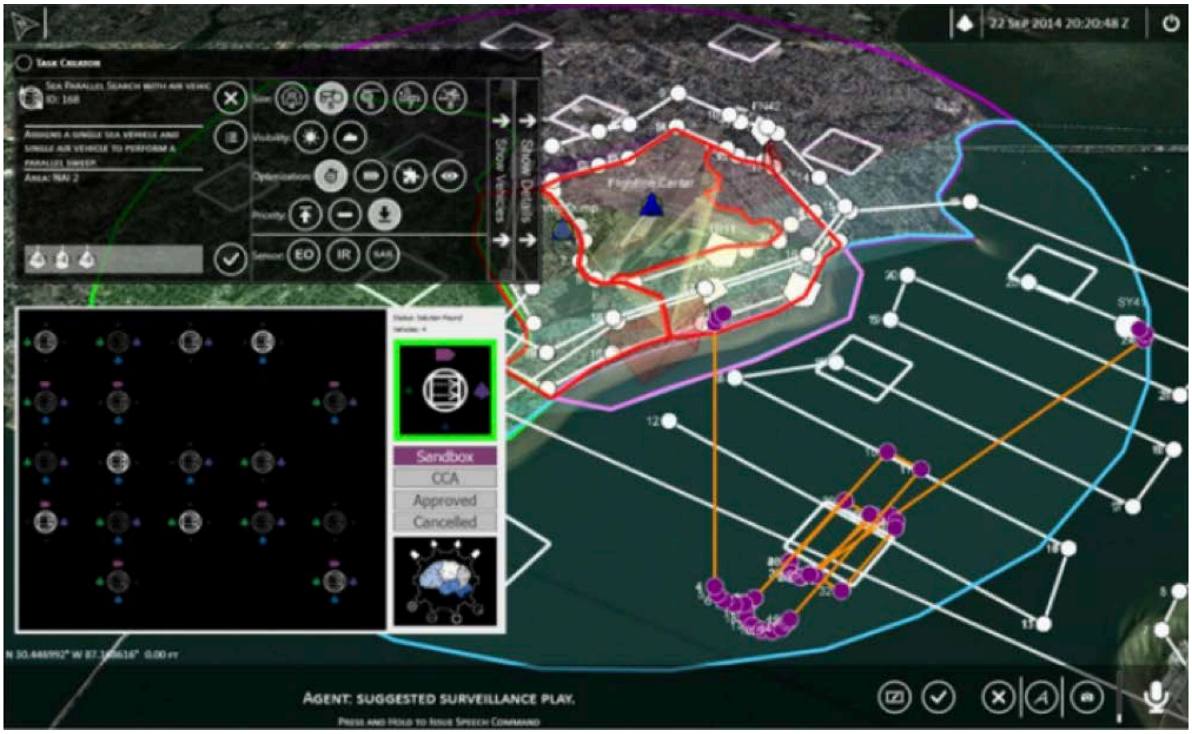

TASK CREATION

SEA PARALLEL SEARCH WITH AIR VEHIC
ID: 168

ASSIGNING A SINGLE SEA VEHICLE AND SINGLE AIR VEHICLE TO PERFORM A PARALLEL SWEEP

AREA: NAI 2

Status: Solution Found
Vehicles: 4

AGENT: SUGGESTED SURVEILLANCE PLAY.

PRESS AND HOLD TO ISSUE SPEECH COMMAND

Figure 2. The IMPACT System.

Within the system (though this can work across system boundaries as well), the information used as part of a planning evolution is tied to the decisions made through the provenance graph. Examples in our case included such situations as the environmental conditions favoring one type of asset over another, or where policies supported the use of one type of asset over another. When these conditions changed during a mission, it was important to reconsider the plan. With the use of provenance graphs we were able to do that. Changes of information tied to the plan resulted in either alerts to the user, or when permitted by the user, could result in AI agents autonomously replanning and changing the missions [3]. In a section below, concerning *explanations*, the manner in which the system would inform the operator of such decisions, will be described.

Opportunities for Learning

Collecting data from instrumented systems and building structures such as provenance graphs, affords the opportunity to develop models for complex analysis. While not yet performed in a military setting, researchers have demonstrated the ability to classify and summarize complex decision strategies from provenance. In [4], a provenance graph kernel is learned and utilized to accomplish these tasks. By observing data over time, an algorithm is able to learn models of different decision strategies. One can then ask the system the question of what decision strategy was used for a particular operation. Similarly, through either this technique or through graph analytics [4, 5], one can develop summaries of large portions of provenance to provide insight into the overall activity of a system or systems.

SUPPORT FOR HYBRID OPERATIONS

Information Across System Boundaries

One of the first instances of PROV-DM serving the information integration of multiple disparate systems, was demonstrated in the Atomic Orchid project in the United Kingdom. Consisting of several different university teams, the project developed command and control capabilities for disaster relief operations [6]. Provenance was used to track the directions of heterogeneous teams of humans and autonomous air vehicles. It was also used for tracking the intelligence analysis using crowd-sourcing. Where system boundaries occur, the opportunity to track decisions without intrusive system changes is presented through a set of provenance instrumentation tools.

A later version of those tools was used for IMPACT [3]. However, the enhanced *Allied* IMPACT (AIM) utilized capabilities outside its boundaries developed by other nations [7, 8]. One such boundary was between IMPACT's plan monitoring capability and a policy tracking capability (COMPACT) developed by the United Kingdom's Defence Science and Technology Laboratory. COMPACT was able to provide predictions of policy violations in a battlespace where policies might change and where systems might not perform exactly as planned. When COMPACT altered information concerning policy compliance, the AIM plan monitor could be re-evaluate the performance of the vehicle relative to the new mission parameters. Because it was permitted by the operator to do so, it would automatically initiate a new planning process, and further utilize a task manager [9], to cause changes in vehicle tasking. All of this would be further tracked in provenance allowing visibility into what agents and processes were used to effect changes in the mission.

Explanations Tied to Knowledge

The IMPACT experiments demonstrated that in a complex environment, balanced use of autonomy within the command and control system allowed the human to manage the scenario and maintain situational awareness [3, 9]. In order for that to work, the human decision-makers needed to be informed when the autonomy made decisions[1]. Speech was chosen as the preferred method for this communication. The system was acting in some ways as a teammate and this was a natural mode of communication for the scope of the decisions being made (e.g., deciding to add a resource to a mission due to a deficiency in a plan currently being executed).

To affect this, we utilized the fact that the provenance trace formed a graph, and that the Australian Defence Science and Technology Group had developed a speech generation capability based on Rhetorical Structure Theory which also utilized graphs [13]. Rhetorical Structure Graphs provide a knowledge-graph like capability, but where nodes are associated with multimedia elements. The capability can actually support providing a multimedia presentation which was utilized in the exercises, but which we did not fully utilize through the translation of provenance. We did however generate speech using it. By translating provenance into this richer graph structure and utilizing some of the knowledge associated with the richer graphs, the system was able to announce using speech the decisions that were made and the reasons for those decisions referencing the events and information that drove them. The knowledge in the graph could be used to provide as much background as needed through connections to doctrine or policy that had nodes and relations available through the graph.

[1] The operator was able to control what decisions the autonomy was allowed to make and under what conditions utilizing a construct known as a working-agreement.

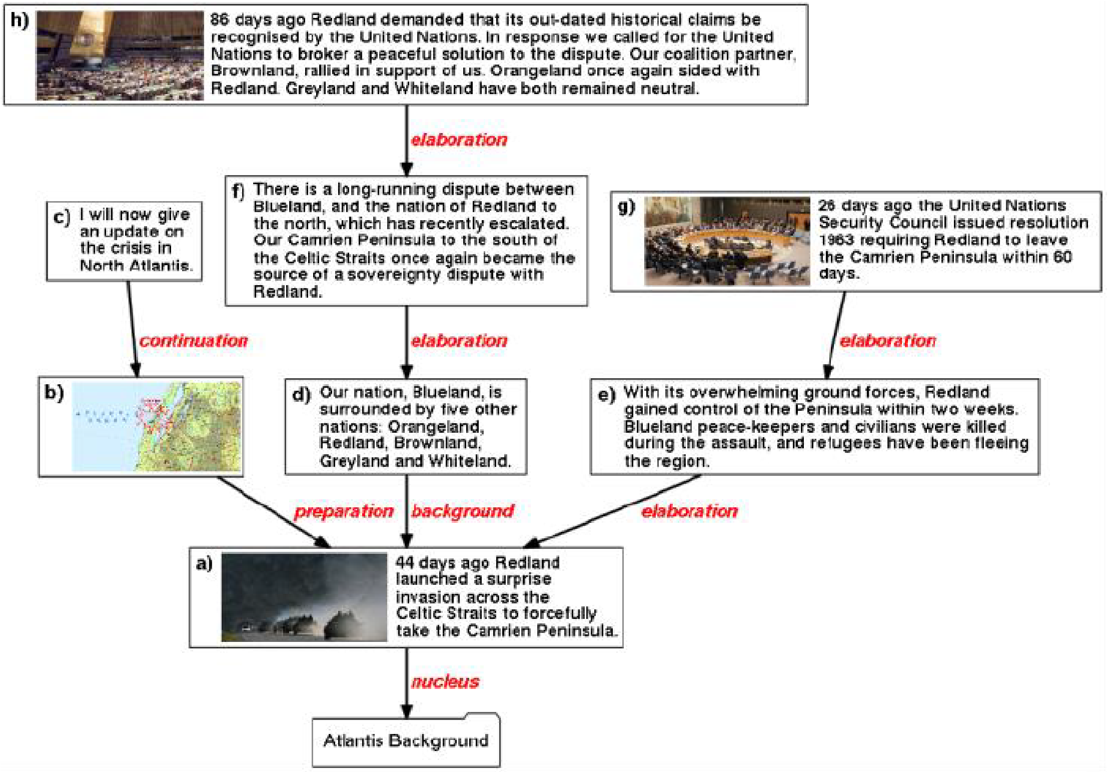
h) 86 days ago Redland demanded that its out-dated historical claims be recognised by the United Nations. In response we called for the United Nations to broker a peaceful solution to the dispute. Our coalition partner, Brownland, rallied in support of us. Orangeland once again sided with Redland. Greyland and Whiteland have both remained neutral.
elaboration
c) I will now give an update on the crisis in North Atlantis.
f) There is a long-running dispute between Blueland, and the nation of Redland to the north, which has recently escalated. Our Camrien Peninsula to the south of the Celtic Straits once again became the source of a sovereignty dispute with Redland.
g) 26 days ago the United Nations Security Council issued resolution 1963 requiring Redland to leave the Camrien Peninsula within 60 days.
continuation
elaboration
elaboration
b)
d) Our nation, Blueland, is surrounded by five other nations: Orangeland, Redland, Brownland, Greyland and Whiteland.
e) With its overwhelming ground forces, Redland gained control of the Peninsula within two weeks. Blueland peace-keepers and civilians were killed during the assault, and refugees have been fleeing the region.
preparation
background
elaboration
a) 44 days ago Redland launched a surprise invasion across the Celtic Straits to forcefully take the Camrien Peninsula.
nucleus
Atlantis Background

Figure 3. Rhetorical relations assigned to multimedia elements. [12]

Coalition Information

Current research is being conducted on developing provenance information in a distributed network under intermittent communications. The canonical scenario is a set of autonomous systems is operating along with a command and control node that has a human operator. At each time step, the systems might be able to communicate with a different subset of the participants. When they can communicate, they transmit their provenance graphs. This allows them all to know what decisions each have made.

Each system has a planner, and through reconstructing the decision environment of a peer can verify that a decision would be appropriate. This can be utilized to evaluate whether or not a member of the network has been corrupted or an attack has occurred resulting in a bad actor being present.

Experiments have been performed utilizing block chain algorithms [14] and algorithms specifically designed for corroboration of provenance [15]. With these approaches, the participants of the network vote on reasonableness of the change utilizing their planning algorithms and their trace of the provenance to decide if a change should be allowed in the provenance as suggested by a participant. This is work in progress and only preliminary results have been achieved.

SUMMARY

Hybrid operations by their very nature will involve a diverse set of systems and information. Humans will be making decisions throughout the network as will AI agents. At the very least this implies a critical task of tracking the decisions made and the contexts in which they were decided. Fundamentally, that is what the provenance graph does. It represents the complex array of decisions and the information environment, within and across systems. But there may be multiple ways to record provenance. This paper argues that the PROV-DM graph structure provides the opportunity for a rich set of capabilities to be developed.

From a forensic point of view, the information is available utilizing the provenance graph. One can query to learn who decided what, using what information. Was it a human or an AI agent that made the decision? Who provide the authorization either for the action or the autonomous behavior?

Further, the provenance graph provides opportunities for capabilities during operations. One of the most interesting is the ability to explain to humans what the autonomy has done and why. The linkage between the provenance graph and rhetorical structure graphs, provides a mapping from observation to knowledge that can come from doctrine and from learning. Using rhetorical structure theory to then generate natural language provides the information during operations in a form that is obtainable for humans.

Not only can provenance track the decisions made by AI that might have models developed through machine learning, but machine learning can be used on the provenance itself. Insight can be gained by classifying decision processes, summarizing the graphs, and extracting metrics. These can deepen our understanding of how a dynamic complex organization of agents, human and artificial, are performing.

Finally, we have the possibility of utilizing provenance to aid in the construction of more global information developed in distributed intermittent communications settings. Utilizing corroboration algorithms, we may be able to identify faults or attacks on a dynamic network of forces and systems that will be required for hybrid warfare.

REFERENCES

[1] Moreau, L., Groth, P. (2013). Provenance: an introduction to PROV. Synthesis Lectures on the Semantic Web: Theory and Technology. Morgan & Claypool.

[2] US ACM. (2017). Statement on algorithmic transparency and accountability. 9(2): 1-2.

[3] Behymer, K., Rothwell, C., Ruff, H., Patzek, M., Calhoun, G., Draper, M., Douglass, S., Kingston, D., Lange, D. (2017). Initial evaluation of the intelligent multi-uxv planner with adaptive collaborative/control technologies (IMPACT). Tech Report, Infoscitex Corp., Beavercreek.

[4] Marzagão, D., Huynh, T., Helal, A., and Moreau, L. (2020), Provenance Graph Kernel. arXiv:2010.10343v1, 20 October 2020.

[5] Huynh, T., Ebden, M., Fischer, J., Roberts, S., and Moreau, L. (2018). Provenance Network Analytics. *Data Mining and Knowledge Discovery*, Feb 2018. ISSN 1384-5810. doi: 10.1007/ s10618-017-0549-3. URL http://link.springer.com/10.1007/s10618-017-0549-3.

[6] Ramchurn, S. D., Huynh, T. D., Wu, F., Ikuno, Y., Flann, J., Moreau, L., Fischer, J. E., et al. (2016). A

disaster response system based on human-agent collectives. Journal of Artificial Intelligence Research, vol. 57, pp. 661-708.

[7] Lucero, C., Huynh, T.D., Lange, D.S., Moreau, L., “Modelling provenance of decisions within the human autonomy team”, NATO STO-MP-HFM-300 Symposium on HAT, 2019.

[8] Coronado, B., and Lange, D., “Autonomic Plan Monitoring”, NATO STO-MP-HFM-300 Symposium on HAT, 2019.

[9] Gustafson, E., Coronado, B., and Lange, D., “Discretizing and managing the task environment”, NATO STO-MP-HFM-300 Symposium on HAT, 2019.

[10] Lucero, C., Coronado, B., Hui, O., Lange, D. S. (2018). Exploring explainable artificial intelligence and autonomy through provenance. IJCAI-ECAI 2nd Workshop on eXplainable Artificial Intelligence (XAI), pp. 85-89.

[11] Lucero, C., Izumigawa, C., Frederiksen, K., Nans, L., Iden, R., Lange, D. S. (2020). Human-autonomy teaming and explainable AI capabilities in RTS games. Engineering Psychology and Cognitive Ergonomics. Cognition and Design. Lecture Notes in Computer Science, vol. 12187, pp. 161-171, Springer, Cham.

[12] Wark, S., Nowina-Krowicki, M., Lucero, C., Lange, D. (2018). But why? Generating narratives using provenance. OzCHI 2018 Workshop on Interaction Design for Explainable AI, pp. 5-8

[13] Wark, S., Nowina-Krowicki, M. (2016). Generating multimedia narrative for virtual humans. Proceedings of the Australasian Computer Science Week Multiconference, pp. 1-10.

[14] Ferdous, M., Chowdhury, M., and Hoque, M. (2020). Blockchain Consensus Algorithms: A Survey. arXiv:2001.07091v2 [cs.DC] 7 Feb 2020.

[15] Barakat, L., Taylor, P., Griffiths, N., and Miles, S. (2017). 9th USENIX Workshop on the Theory and Practice of Provenance, Seattle, Washington, 23 June 2017.