Title page

Authors: Steve North

Corresponding author:
Steve North
Contact address:
2 Height Top Cottages
Badger Lane
Blackshaw Head
Hebden Bridge
West Yorkshire HX7 7JW
United Kingdom
Telephone: +44 (0)1422-847-254
Email: steve@stevenorth.com

Suggested title: '**Visualizing decision-making behaviours in agent-based autonomous spacecraft**'

Abstract

The authors will report initial progress on the PIAudit project as a Research Resident Associate Program. The objective of this research is to prototype a tool for visualizing decision-making behaviours in autonomous spacecraft. This visualization will serve as an information source for human analysts. The current visualization prototype for PIAudit combines traditional Decision Trees with Weights of Evidence.

1. Introduction

As NASA systems introduce higher levels of autonomy ([1][2][3]), it is increasingly important that such systems have the functionality to justify automated decisions. This paper describes one attempt to analyze and visually represent decision-making behaviors. The object of the PIAudit (Plan Audit) research project is to prototype an approach to visualizing decision information. Specifically, PIAudit visualizes the decisions made by an autonomous spacecraft's onboard software agents.

For the purposes of prototyping, PIAudit uses NASA GSFC's Autonomous Nano Technology Swarm (ANTS) concept as a source of decision data. ANTS has a possible NASA implementation date of 2020. To summarize, ANTS consists of one thousand pico spacecraft (with a mass of less than one kilogram each) flying from Earth orbit to the asteroid belt, using solar sails. [4][5] In the prototype, ANTS will be simulated as a computer-based multi-agent system, communicating decision data to PIAudit.

In the second section of the paper, an introduction is provided to the use of Information Visualization technologies at NASA Goddard Space Flight Center in the Advanced Architectures and Automation Branch. The next section provides a

brief introduction to the 'state of the art' in analyzing decisions made by autonomous spacecraft. The analysis and representation of all types of decision-making is introduced in Section 4. The PIAudit architecture and experimental context is then introduced.

2. Information Visualization at NASA Goddard Space Flight Center in the Advanced Architectures and Automation Branch

Before considering the specific case of visualizing decisions, this section will briefly describe the use of Information Visualization in the Advanced Architectures and Automation Branch (http://aaaprod.gsfc.nasa.gov) at NASA Goddard Space Flight Center.

NASA Goddard Space Flight Center has the main mission of conducting science observations of both the Earth and Space largely from satellites in low Earth orbit. This wide variety of possible missions and spacecrafts generates many types of telemetry data, which are visualized in various ways. Some projects at NASA Goddard which have used data visualization to represent satellite data are VISAGE (Visual Analysis Graphical Environment - http://invision.gsfc.nasa.gov/avatar/projects/visage/index_oldversion.html), IRC (Instrument Remote Control - http://pioneer.gsfc.nasa.gov/public/irc/) and REACH (Realtime Evaluation and Analysis of Consolidated Health - http://invision.gsfc.nasa.gov/avatar/projects/reach/).

Each of these projects has examined both discrete and continuous spacecraft data variables using custom visualizations for a variety of spacecrafts. However, few spacecrafts up to the present time (2003) have included much in the way of onboard autonomous systems. It is *critical* to mission success that the decision processes of future autonomous systems be viewable easily by human

operators. For this reason it is essential that accurate, data-rich, easily readable visualizations be developed to support these missions. In this way autonomous satellite decisions may become trusted by human ground system operators so that more autonomy may be infused into future missions which will expedite science data return and lower mission cost.

As we have explained above, the use of data visualization in the representation of decision data related to intelligent, autonomous spacecraft is a critical, but largely unexplored and emerging area of research at NASA. However, non-visual analysis techniques have been studied. These approaches are considered next.

3. Existing non-visual techniques for auditing decision-making behaviours in autonomous spacecraft

The purpose of this section is to briefly describe contemporary efforts to report the behaviour of autonomous spacecraft. This overview is provided to contextualize proposals in later sections for the application of Information Visualization.

NASA has been operating missions for many years that include low-level elements of autonomy. Such missions encounter the problem of retrospectively analyzing the behavioral outcomes of autonomous decision-making. For example, NASA's 1997 Mars Sojourner microrover featured limited autonomous behaviour. The literature (for example Mathies et al. [6] and Laubach et al. [7]) suggests that Sojourner did not log supporting evidence for its autonomous decisions and that such an approach was not available to its engineers. Retrospective analysis of Sojourner's self-navigation seems to have relied on studying actual paths taken to reach 'waypoints', in conjunction with the 3D terrain maps generated by the rover's stereo cameras.

Post-Sojourner, work concerned with planetary rover control has tended to focus on improving autonomous decision-making. However, little attention seems to have been paid to the remote human operator's ability to verify autonomous decisions. Given the increasing complexity of demands placed on system autonomy, it is surprising that research into legacy diagnosis and analysis has

not kept pace. Where diagnostic analysis is applied, it frequently seems to focus on the study of the result of the decision, rather than its supporting beliefs.

Richard Washington and collaborators at NASA Ames Research Center seem to have been authoritative in progressing autonomous control. Washington's work on 'decision-theoretic' planning is particularly important. [8][9][10][11][12] The term decision-theoretic refers to decision theory as first described by Luce & Raiffa. [13] This provides a framework for weighing the strengths and weaknesses of a particular course of action. Given a probability distribution over the possible outcomes of an action (in any state), it is possible to grade potential plans according to likelihood of success.

Fault protection is one area where the need for validation of autonomy has been investigated. A concept known as, 'Beacon Mode Operations' has been developed, wherein an autonomous spacecraft sends one of several tone signals to request ground action. [14] Beacon Mode Operations include:

> "onboard engineering-data summarization in an ongoing fashion, so that when an emergency signal arrives from the spacecraft, it is quickly followed- once a full communications link is established- by an anomaly report, including context and completed analysis, to bootstrap the ground-based troubleshooting effort". [15]

NASA's Deep Space One (DS1) mission utilized 'behaviour auditors' to monitor the runtime execution of its autonomy software. [16] A fault protection component within the spacecraft avionics logged all responses to faults for communication to the ground, as part of a Beacon Mode system. Bernard et al. state that, "the

encoded path and ancillary sampled variables constitute...event records to provide full accountability of the rationale behind every state transition for every fault-response execution". [17]

Despite certain conceptual similarities between DS1's behavior auditors and proposals for PlAudit, there are several distinctions. Primarily, DS1 only logged statecharts for the rationale behind responses to faults. In addition to having a different probabilistic model, PlAudit focuses on *visualizing* the evidence behind decisions. PlAudit also logs all decisions, not just those made in response to fault detection.

To summarize, little or no work seems to have been done on archiving and analyzing the justification or evidence behind the decisions made by autonomous spacecraft. Therefore, PlAudit is required to provide an underlying approach to logging decision evidence, in addition to visualizing it. The next section considers techniques for visualizing decisions and evidence in non space-related domains. The objective being to identify approaches that might be generalized.

4. Traditional approaches to analysis and visualization of decision-making

One of the most common methods for representing decision-making is the 'decision tree'. This term has various definitions depending on the domain of application. For the purposes of PIAudit, it can be assumed that the MIT Encyclopedia of Cognitive Science definition applies: "a graphical representation of a procedure for classifying or evaluating an item of interest". [18] When used to represent a temporal sequence of decisions (as used here), the tree can be considered to be 'classifying' rather than 'evaluating' items. This contrasts with the more rigorous usage of the term in the field of information theory, that describes a decision tree as consisting of nodes and branches determined by their information content. [19]
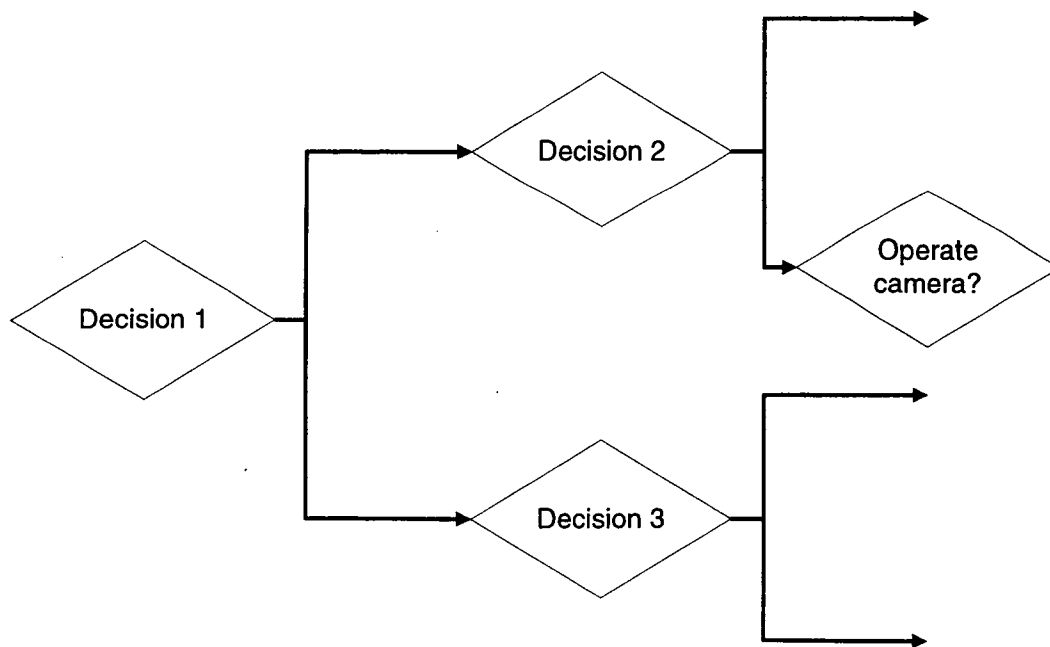

Figure 1 Example decision tree

Figure 1 illustrates a simple decision tree, showing all possible routes. From the

perspective of PIAudit's analysis of *actual* decisions taken, the visualization

interface only needs to display the chosen path, not all branches.

An alternative graphing view allows the generation of a Belief Network. This is:

> "a data structure (a directed acyclic graph) that is used to represent
> dependence between variables. Each variable has a corresponding node
> with a conditional probability table defining the relationships between
> parent nodes. The primary use of Bayesian networks is to use probability
> theory to reason with uncertainty". [20]

A belief network is excellent at demonstrating interdependence between

variables, but it does not provide any indication of the relative importance of each

variable to the outcome. In the case of PIAudit, the outcome would correspond to

making a decision. The variables would represent the evidence considered by an

agent when evaluating a decision.

Good [21] suggests 'Weights of Evidence' as an approach to quantifying 'relative

importance to outcome' or 'explanatory importance' (as it is more usually

described). Good's algorithms make use of Bayes Theorem, which states that:

> "we can compute conditional probability that event Y will occur given that
> event X already occurred, providing we know the prior probability that X
> and Y could happen, and the conditional probability that X will occur when
> we know that Y has already occurred". [22]

The graphical explanation work of Almond is seminal in combining belief

networks with weights of evidence. [23][24] Almond and collaborators present

methods for visualizing probabilistic 'evidence flows' in belief networks, thereby

enabling belief networks to explain their behavior. This delivers a hierarchy of explanations, ranging from simple colorings to detailed displays.

Figure 2 shows a simple belief network with node coloring representing weights of evidence. Without knowing the exact probabilities, it is still possible to deduce that the camera being online and focused provides strong positive evidence for operating the camera. This evidence outweighs the weak negative evidence provided by the lighting conditions being poor. Of the two unknown variables, the agent should attempt to ascertain whether the target object is framed. This variable has the greatest estimated (or potential) weight of evidence.
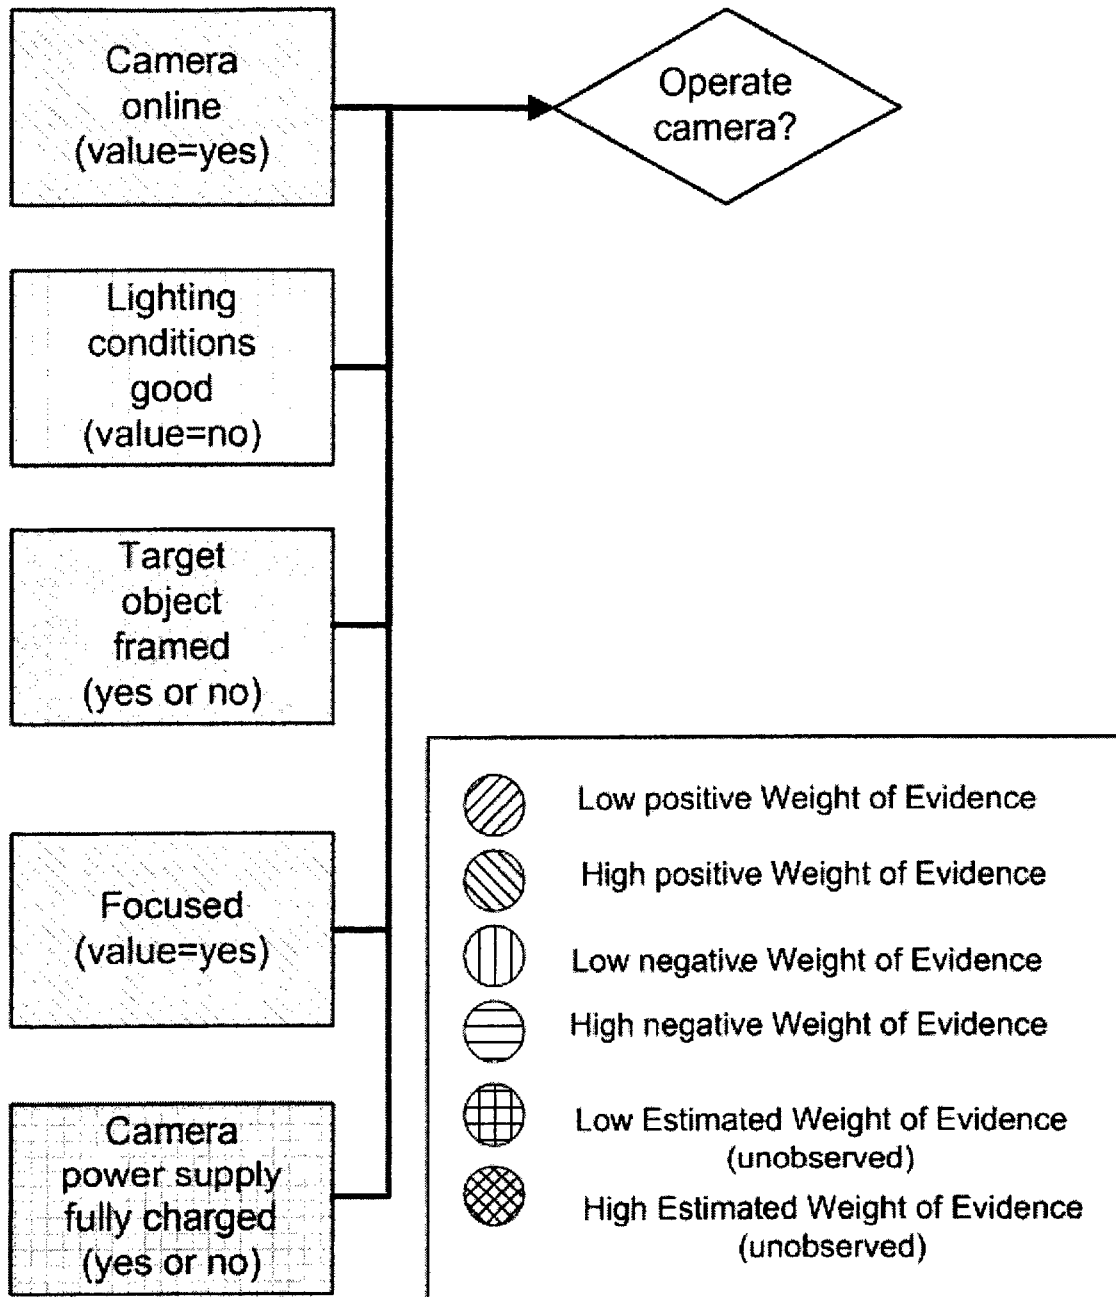
Figure 2 Example Belief Network with Weights of Evidence

Weights of evidence are not the only methodology devised for quantifying degree

of belief. For example, Xu et al. [25] propose the Transferable Belief Model

(TBM). This system performs evidential reasoning and decision-making. It

achieves this by integrating an evidential system for belief function propagation and a valuation-based system for Bayesian decision analysis.

5. The PIAudit system architecture

The last two sections have clarified two issues. Firstly, the need for a tool to visualize the evidence behind the decisions made by autonomous spacecraft. Secondly, the availability of decision visualization techniques in other fields (decision trees, belief networks etc.). This section describes how the PIAudit prototype attempts to unite the identified problem with a possible solution.

Figure 3 illustrates how PIAudit operates as an intermediary between the decision-making of one or more agents and a human operator. As previously stated, the PIAudit prototype will exchange decision data with a multi-agent system, representing a satellite constellation mission.
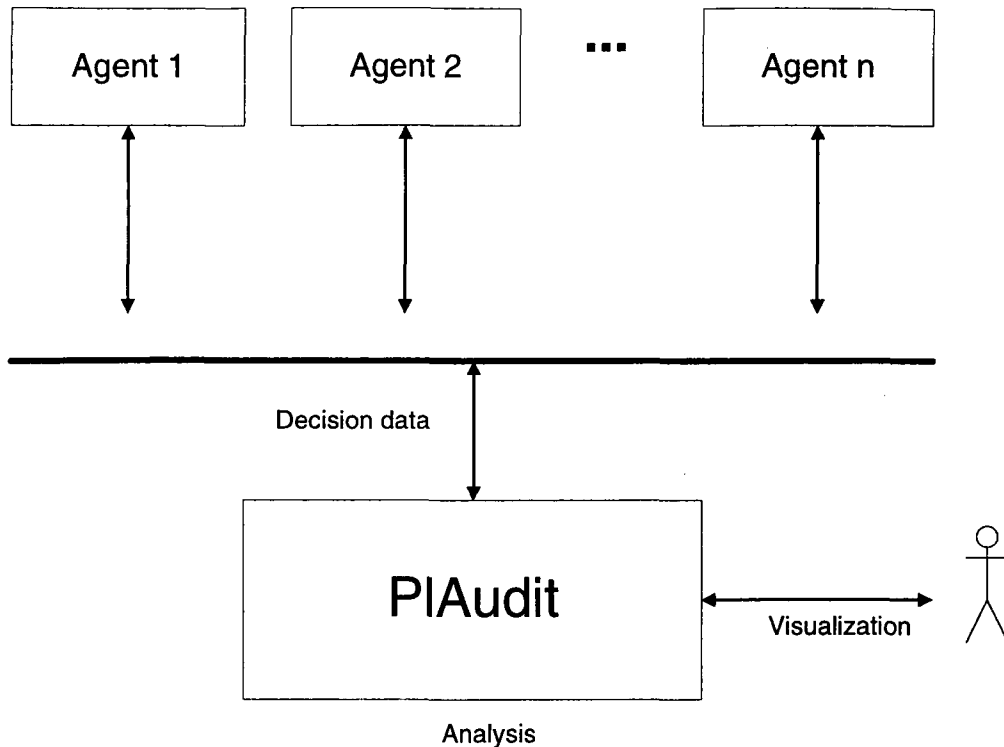
Figure 3 The PIAudit architecture

The last section introduced decision trees, belief networks and weights of evidence. The PIAudit methodology combines decision trees and weights of evidence. Each decision node in the tree is supported by a chain of supporting evidence, explaining a system's rationale for each choice. The process for generating a visual structure from decision data is as follows:

➢ Identify a decision and its possible values. For the purposes of mathematical simplicity, the initial PIAudit prototype assumes Boolean true or false decisions.

➢ Identify the indicants that impact this decision. Indicants are the variables (again valued true or false) that may inform the selection of choices for

this decision. The decision-maker will control some indicants, but not necessarily all. Indicants can be thought of as encompassing both precepts (states perceived by but not under the direct control of the decision-maker) and internal states, which are decision-maker controlled.

➢ For each indicant, identify the specific nature of the Boolean values (i.e. true/false, empty/full, <1 / >1 etc.).

➢ Define the 'beliefs' or 'expert knowledge' that determine how the various values of indicants should influence choice for this decision. For the values of each indicant specify the probability that this value should suggest the selection of a particular decision choice. Repeat for each binary indicant:

$$Pr\ (\ Indicant=0\ |\ Choice=0\ )$$
$$Pr\ (\ Indicant=0\ |\ Choice=1\ )$$
$$Pr\ (\ Indicant=1\ |\ Choice=0\ )$$
$$Pr\ (\ Indicant=1\ |\ Choice=1\ )$$

➢ For the values of each indicant, calculate the Actual Weight of Evidence (AWOE). Based on the belief probabilities, this represents the amount of evidence for (or against) selecting a particular decision choice. It is the evidence present if THIS indicant is in THIS state:

$$AWOE\ (\ Indicant = 0\ ) = \quad Log\ \frac{Pr\ (\ Indicant=0\ |\ Choice=1\ )}{Pr\ (\ Indicant=0\ |\ Choice=0\ )}$$

$$AWOE\ (\ Indicant = 1\ ) = \quad Log\ \frac{Pr\ (\ Indicant=1\ |\ Choice=1\ )}{Pr\ (\ Indicant=1\ |\ Choice=0\ )}$$

16

➤ Calculate the Expected Weights of Evidence (EWOE) for each indicant. This value can be thought of as a 'measure of the information content of a future finding'. Another way of looking at this is that it represents the potential AWOE that might better inform making the decision, where a particular indicant value is unknown. EWOE does not tell you whether the evidence is positive or negative, just the strength of the evidence. When trying to decide which unobserved indicants to pursue, EWOE provides a basis for evidence selection. The algorithm is as follows:

EWOE (Indicant) =
Σ (AWOE (Indicant=value) * Pr (Indicant=value | Choice = 1))

➤ In the case of our binary indicants, this can be read as:

EWOE (Indicant) =
(AWOE (Indicant=0) * Pr (Indicant=0 | Choice = 1))
+
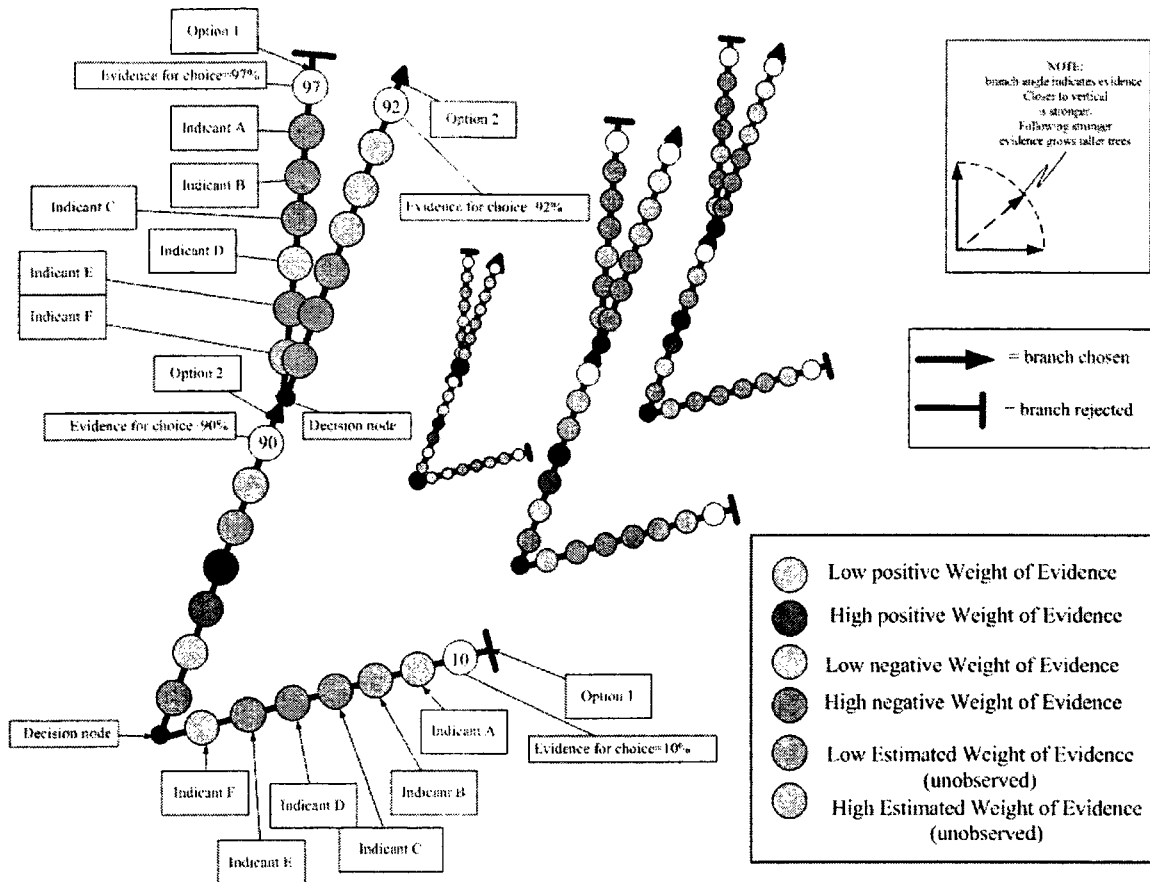(AWOE (Indicant=1) * Pr (Indicant=1 | Choice = 1))

Figure 4 Preliminary design for PIAudit visualization

Figure 4 illustrates the authors' preliminary design for an Information

Visualization, utilising Weights of Evidence and Decision Trees. Card et al. define

Information Visualization as: "the use of computer-supported, interactive, visual

representations of abstract data to amplify cognition". [26] PIAudit uses

Information Visualization to explore the non-physical data sets of decision-

making. PIAudit's Information Visualization is realised using eXtensible 3D (X3D).

[27]

The final design of PIAudit's visualization is expected to use the third spatial

dimension to represent relationships between separate agents (ANTS in this

context). Each tree in the 'forest' displays the decisions taken by individual agents. PIAudit is implemented in the Java programming language. The ANTS simulation is also Java powered, using the MadKit multi-agent platform. [28] Figure 5 shows the PIAudit application's communication interfaces.
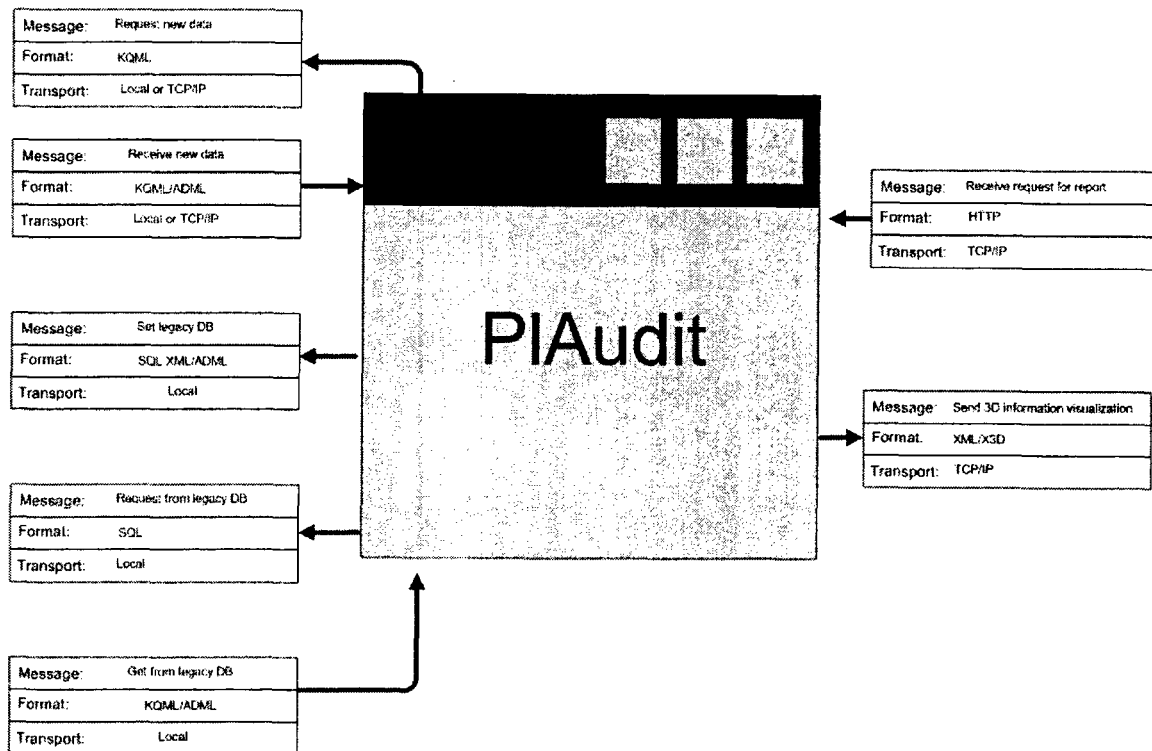

Figure 5 PIAudit interfaces

Figure 6 PIAudit experimental context

Figure 6 shows a proposed experimental setup for evaluating the PIAudit
prototype. Three separate computers are used to represent the physically
separated 'Space', 'Ground' and 'Operator' elements. The ground-based PIAudit
component requests and receives decision-making data from members of the
space-based, autonomous satellite constellation. PIAudit archives decision-data
in a local database store. On an Internet (TCP/IP) transported request from the
human operator, PIAudit generates and returns a visualization. This is sourced
both from the data store and the simulated satellites, as appropriate.

6. Conclusions

This paper has reported the current status of NASA's PIAudit project. A brief introduction to the use of Information Visualization at NASA Goddard Space Flight Center in the Advanced Architectures and Automation Branch was provided in Section 2. In the third section of this paper, it was noted that most existing work on spacecraft decision auditing focuses on post-decision *consequences* rather than pre-decision *evidence*. Recent work on Beacon Mode Operations has built on this consequences-only model, by requiring a system to explain the rationale (or evidence) behind its decisions. A review of the literature did not identify any examples in which visualization techniques had been applied to this problem.

As there appeared to be no current literature concerning the visualization of decision analysis in the field of spacecraft autonomy, the fourth section of this paper considered visual representations of decisions in other domains. PIAudit develops Almond's work on combining belief networks with weights of evidence. With PIAudit, weights of evidence are combined with decisions trees and used to analyse the behaviour of autonomous spacecraft. The fifth section of this paper introduced the PIAudit architecture and the planned experimental context for evaluating its prototype. Further work will focus on further visualization designs and developing an XML dialect for communicating decision data between the ANTS and PIAudit. [29]

7. References

1. Doyle R (1998a). Spacecraft autonomy and the missions of exploration. *IEEE Intelligent Systems.* September/October: 36-44.

2. Doyle R (1998b). What does autonomy capability really mean for JPL and NASA? A talk with Robert Rasmussen. *IEEE Intelligent Systems.* September/October: 76-77.

3. Chweh C (1998). Autonomy in space. *IEEE Intelligent Systems.* September/October: 78-80.

4. Curtis S A, Mica J, Nuth J, Marr G, Rilee M & Bhat M (2000). ANTS (Autonomous Nano-Technology Swarm): an artificial intelligence approach to asteroid belt resource exploration in *Proceedings of International Astronautical Federation, 51st Congress.*

5. NASA (2002c). *ANTS home page* [WWW document] http://ants.gsfc.nasa.gov/ (accessed October 2002).

6. Mathies L, Gat E, Harrison R, Wilcox B, Volpe R & Litwin T (1995). Mars microrover navigation: performance evaluation and enhancement. *Autonomous Robots, Special Issue on Autonomous Vehicles for Planetary Exploration.* 2(4): 291-311.

7. Laubach S L & Burdick J W (1999). An Autonomous Sensor-Based Path-Planner for Planetary Microrovers'. *Proceedings of IEEE International Conference on Robotics and Automation.*

8. Washington R, Golden K & Bresina J (1999). Plan Execution, Monitoring, and Adaptation for Planetary Rovers. *Proceedings of the IJCAI-99 Workshop, Scheduling and Planning meet Real-time Monitoring in a Dynamic and Uncertain World.*

9. Bresina J & Washington R (2000). Expected Utility Distributions for Flexible, Contingent Execution *Proceedings of the AAAI-2000 Workshop: Representation Issues for Real-World Planning Systems.*

10. Bresina J & Washington R (2001). Robustness via Run-time Adaptation of Contingent Plans. *Proceedings of the AAAI-2001 Spring Syposium: Robust Autonomy.* Stanford.

11. Zilberstein S, Washington R, Bernstein D & Mouaddib A (2001). Decision-theoretic control of planetary rovers. *Proceedings of the Dagstuhl Workshop on Plan-based Control of Robotic Agents.* Wadern. Germany. 2001.

12. Bernstein D, Zilberstein S, Washington R & Bresina J (2001). Planetary rover control as a markov decision process. *Proceedings of the AAAI Spring Symposium: Game Theoretic and Decision Theoretic Agents.*

13. Luce R & Raiffa H (1957). *Games and decisions: introduction and critical survey.* USA: Wiley.

14. Bernard D, Doyle R, Riedel E, Rouquette N, Wyatt J, Lowry M & Nayak P (1999). Autonomy and software technology on NASA's Deep Space One. *IEEE Intelligent Systems.* May/June: 10-15.

15. Doyle R (1998). Spacecraft autonomy and the missions of exploration. *IEEE Intelligent Systems.* September/October: 36-44.

16. Bernard D, Doyle R, Riedel E, Rouquette N, Wyatt J, Lowry M & Nayak P (1999). Autonomy and software technology on NASA's Deep Space One. *IEEE Intelligent Systems.* May/June: 10-15.

17. Bernard D, Doyle R, Riedel E, Rouquette N, Wyatt J, Lowry M & Nayak P (1999). Autonomy and software technology on NASA's Deep Space One. *IEEE Intelligent Systems.* May/June: 10-15.

18. MIT (2002). MIT Encyclopedia of Cognitive Science [WWW document] http://cognet.mit.edu/MITECS/Entry/utgoff.html (accessed October 2002).

19. Bigus P & Bigus J (2001). *Constructing Intelligent Agents Using Java.* USA: Wiley.

20. Bigus P & Bigus J (2001). *Constructing Intelligent Agents Using Java.* USA: Wiley.

21. Good I J (1950). *Probability and the Weighing of Evidence.* Hafner: New York.

22. Bigus P & Bigus J (2001). *Constructing Intelligent Agents Using Java.* USA: Wiley.

23. Almond R (1994). *Capturing Reliability Knowledge in GRAPHICAL-BELIEF.* StatSci Research Report 31.

24. Madigan D, Mosurski K & Almond R (1997). Graphical explanation in belief networks. *Journal of Computational and Graphical Statistics.* 6 (2): 160-181.

25. Xu H, Hsia Y & Smets P (1995). Transferable Belief Model for Decision Making in Valuation-Based Systems. *IEEE Transactions on Systems, Man and Cybernetics.*

26. Card S K, Macklinlay J D & Shneiderman B (1999). *Readings in Information Visualization- using vision to think.* USA: Morgan Kaufmann.

27. Web3D Consortium (2001). X3D - Extensible 3D: New-Generation Open Web3D Standard [WWW document] http://www.web3d.org/x3d/ (accessed January 2003).

28. MadKit (2003). Home Page [WWW document] http://www.madkit.org/ (accessed January 2003).

29. NASA (2003). PlAudit Home Page [WWW document] http://agents.gsfc.nasa.gov/research/plaudit/index.html (accessed January 2003).