

AFRL-RI-RS-TR-2009-102
Final Technical Report
April 2009



PROACTIVE INTELLIGENCE (PAINT) SIMULATED EXPLORATION OF EXECUTABLE DESIGN STRATEGIES (SEEDS)

Techteam Government Solutions, Inc.

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

STINFO COPY

**AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE
ROME RESEARCH SITE
ROME, NEW YORK**

NOTICE AND SIGNATURE PAGE

Using Government drawings, specifications, or other data included in this document for any purpose other than Government procurement does not in any way obligate the U.S. Government. The fact that the Government formulated or supplied the drawings, specifications, or other data does not license the holder or any other person or corporation; or convey any rights or permission to manufacture, use, or sell any patented invention that may relate to them.

This report was cleared for public release by the 88th ABW, Wright-Patterson AFB Public Affairs Office and is available to the general public, including foreign nationals. Copies may be obtained from the Defense Technical Information Center (DTIC) (<http://www.dtic.mil>).

AFRL-RI-RS-TR-2009-102 HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION STATEMENT.

FOR THE DIRECTOR:

/s/

NANCY A. ROBERTS
Work Unit Manager

/s/

JOSEPH CAMERA, Chief
Information & Intelligence Exploitation Division
Information Directorate

This report is published in the interest of scientific and technical information exchange, and its publication does not constitute the Government's approval or disapproval of its ideas or findings.

REPORT DOCUMENTATION PAGE*Form Approved*
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Service, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington, DC 20503.

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.**1. REPORT DATE (DD-MM-YYYY)**

APR 09

2. REPORT TYPE

Final

3. DATES COVERED (From - To)

Dec 07 – Jan 09

4. TITLE AND SUBTITLE

PROACTIVE INTELLIGENCE (PAINT) SIMULATED EXPLORATION OF EXECUTABLE DESIGN STRATEGIES (SEEDS)

5a. CONTRACT NUMBER

FA8750-08-C-0012

5b. GRANT NUMBER

N/A

5c. PROGRAM ELEMENT NUMBER

N/A

6. AUTHOR(S)

Rafael Alonso, Sven Brueckner, Diane Yang, and Andrew Yinger

5d. PROJECT NUMBER

PAIN

5e. TASK NUMBER

00

5f. WORK UNIT NUMBER

14

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)Techteam Government Solutions, Inc.
3863 Centerview Drive, Suite 150
Chantilly, VA 20151-3287**8. PERFORMING ORGANIZATION
REPORT NUMBER**

N/A

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)AFRL/RIED
525 Brooks Rd.
Rome NY 13441-4505**10. SPONSOR/MONITOR'S ACRONYM(S)**

N/A

**11. SPONSORING/MONITORING
AGENCY REPORT NUMBER**
AFRL-RI-RS-TR-2009-102**12. DISTRIBUTION AVAILABILITY STATEMENT**

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED. PA# 88ABW-2009-1389

13. SUPPLEMENTARY NOTES**14. ABSTRACT**

The ProActive Intelligence(PAINT) Simulated Exploration of Executable Design Strategies (SEEDS) project first defined the overall systems architecture for the creation of robust, complex, & optimized probes based on a generate & test approach, then specified the architecture & processes within the Possible World Generator, our key component for testing & evaluating probe candidates, & finally, implemented & demonstrated an illustration-of-concept interactive prototype that shows the primary interactions of the key components of the Possible World Generator.

15. SUBJECT TERMS

proactive intelligence, prediction, design strategies, systems architecture, leadership modeling, test probes

16. SECURITY CLASSIFICATION OF:**a. REPORT**
U**b. ABSTRACT**
U**c. THIS PAGE**
U**17. LIMITATION OF
ABSTRACT**

UU

**18. NUMBER
OF PAGES**

28

19a. NAME OF RESPONSIBLE PERSON

Nancy A. Roberts

19b. TELEPHONE NUMBER (Include area code)

N/A

TABLE OF CONTENTS

1	ACCOMPLISHMENTS	1
1.1	PAINT SEEDS ARCHITECTURE	1
1.1.1	<i>Polyagent Emulation Derived from Integrated Models</i>	<i>2</i>
1.1.2	<i>Target Models Driving Polyagents.....</i>	<i>2</i>
1.1.3	<i>Symbolic Reasoning to Generate Probes</i>	<i>3</i>
1.1.4	<i>Adaptive Search over Probe Fitness Landscapes</i>	<i>3</i>
1.2	POSSIBLE WORLD GENERATOR	4
1.2.1	<i>The PAINT SEEDS Pathway Model</i>	<i>7</i>
1.2.2	<i>The PAINT SEEDS Leadership Model</i>	<i>10</i>
1.2.3	<i>Executing the Two Models.....</i>	<i>12</i>
1.3	ILLUSTRATION-OF-CONCEPT DEMONSTRATION	17
1.3.1	<i>Pathway Model.....</i>	<i>17</i>
1.3.2	<i>Leadership/Target Model</i>	<i>18</i>
2	CONCLUSION	21
3	LIST of ACRONYMS.....	22

LIST OF FIGURES

Figure 1. Pathways to Nuclear Technology.....	2
Figure 2. Multiple Search Agents exploring the Space of Possible COA's driven by Solvers and Modifiers.....	3
Figure 3. PAINT SEEDS Architecture for the adaptive search of the space of possible probes.	5
Figure 4. The Possible World Generator's dynamics must be similar to the dynamics of the PAINT Integrated Model.....	5
Figure 5. The Possible World Generator's high-level components match those of the PAINT Integrated Model.....	6
Figure 6. Both the leadership and the social network are graph-based to provide topologies for sub-symbolic agent interactions.	6
Figure 7. Nodes in the pathway model represent necessary technology advancement actions.	7
Figure 8. Resources may be allocated to any and all nodes in the pathway model relative to an estimate of “now” – the current state of the technology development.....	8
Figure 9. Pathway polyagents of competing hypotheses sample the pathway model.....	8
Figure 10. A “Book of Pheromones” captures the alternative futures explored by the pathway ghosts.....	9
Figure 11. The most likely path in the “Book of Pheromones”.....	9
Figure 12. Entities and their social-network relations in the leadership model.....	10
Figure 13. Some entities have means to generate resources.	10
Figure 14. A spreading “activation” process distributes resources to other nodes in the social network.....	11
Figure 15. Entity avatars associated with each node in the leadership model.....	11
Figure 16. Entity avatars contain a target model that determines the allocation of resources available to this entity.	12
Figure 17. Steps 0-1.	13
Figure 18. Steps 2-4.	13
Figure 19. Steps 5-7.	14

Figure 20. Steps 8-11.	15
Figure 21. Steps 12-END.....	16
Figure 22. The Nuclear Technology Development Pathway Model.....	17
Figure 23. The Nuclear Technology Development Leadership and Target Models.....	19

1 ACCOMPLISHMENTS

In the work package executed with the available funding in the Pro Active Intelligence(PAINT) Simulated Exploration of Executable Design Strategies(SEEDS) project we first defined the overall systems architecture for the creation of robust, complex, and optimized probes based on a generate and test approach, then we specified the architecture and processes within the Possible World Generator, our key component for testing and evaluating probe candidates, and finally, we implemented and demonstrated an illustration-of-concept interactive prototype that shows the primary interactions of the key components of the Possible World Generator.

1.1 PAINT SEEDS ARCHITECTURE

In the PAINT SEEDS architecture we define four main components supporting a generate-and-test strategy to search for robust, complex, and optimized probes:

1. A low-dimensional polyagent simulation generating multiple futures concurrently, with the model derived from the high-dimensional models of the other PAINT performers, structured into a leadership model with target models representing the goals of individual leadership actors, and a pathway model representing the constraints of technology development
2. A trajectory evaluation component that evaluates the effect of the execution of a particular probe candidate on the predicted evolution of the system under analysis and our ability to discern or influence the intent and plans of the leadership in terms of the specific outcome of the technology development
3. An adaptive distributed search infrastructure that guides the generation, evaluation, and quantitative modification of probe candidates by exploring the space of possible probes
4. One or more probe generator components that reason about the composition of the model of the system (pathway model and leadership model with targets models) under consideration to derive likely probe candidates.

In the following, we discuss components 1, 2 and 4 in detail. Due to the reduction of scope of our Statement of Work in the funding cut, we had to leave significant portions of the architecture (in particular the integration with other performers) at the planning stage and focused on implementing an illustration-of-concept demonstration (Section 1.3) of the core process flow inside the Possible World Generator. Further refinements and implementation efforts have been deferred to possible future work packages.

1.1.1 POLYAGENT EMULATION DERIVED FROM INTEGRATED MODELS

The Component Predictive Models (CPM) from various PAINT contractors are going to be integrated through the British Aerospace's (BAE's) integration platform to form a coherent and very detailed single-trajectory simulation model of the leadership of a nation state and its alternative pathways to attaining significant capabilities in a dual-use technology area. To design effective probes of the leadership-pathway system, we need a fast emulation of this complex simulation that allows us to explore the performance of many candidate probes and the "bundles" of probable futures that result from the execution (or lack thereof) of an active probe. Models comprising polyagents (complex avatar agents guided by swarms of simplistic behavioral ghost agents) offer such efficient multi-future capability.




FIGURE 1. PATHWAYS TO NUCLEAR TECHNOLOGY.

After extensive discussion with the system integrator team (BAE) as well as the CPM teams, we decided to derive polyagent leadership-pathway models semi-automatically from data provided by and observations obtained from the execution of the complex single pathway model. In these models, we will have polyagents representing entities in the social network of the leadership as well as polyagents representing components of the dual-use technology program progressing on pathway models under the influence of external conditions as well as internal decisions made by the leadership. Figure 1 shows an example for such a pathway model in the context of nuclear technology development.

By itself, the polyagent model would reflect the transition of the state's dual-use technology development program along the known pathways into alternative futures without being influenced by any active probe. From the bundle of alternative futures, we could estimate at what point in the future our Intelligence Community would be able to detect with sufficient certainty, whether the technology was developed for benign or nefarious use. Our polyagent model would also be capable of emulating the execution of an active probe (series of direct or indirect actions on the system at hand) to estimate how much earlier (if at all) we would reach this point and under which side-effects. We turn this information into a search-guiding fitness function for our search architecture discussed below.

1.1.2 TARGET MODELS DRIVING POLYAGENTS


Current research approach consists of modeling the state's leadership along both social network and governmental responsibility dimensions. Base information could be obtained from other program teams, and models will be constructed showing positive/negative relationships and organizational control over segments of the economy and military. Adaptation algorithms will be developed to gauge targets' degree of response to probes as well as possible probe extent. Feedback mechanisms based on documents labeled via the Language Computer Corporation (LCC) Natural Language Processing (NLP) processes.

1.1.3 SYMBOLIC REASONING TO GENERATE PROBES

Our approach to strategy generation works with futures (trajectories) generated by polyagents and is guided by entity modeling of targets (i.e. leaderships). We would first develop a scoring function to assess the goodness (i.e. the extent of being “good” or “bad”) of a particular future based on entity models of targets and analysts. Given a desired end state for a future, i.e. a particular goodness score, we will reason “backwards” to identify key system variables for probing. In other words, we will systematically vary the values of system variables and observe the corresponding changes in the resultant score of the future. We will then identify key variables which cause maximal changes in the goodness score with minimal adjustment to the variables’ values. These key variables become the strategic or binding locations for probing. In this context, a probe is an outside event that effects a change in a system variable. A strategy is a plan of binding one or more probes to their respective key variables to achieve a desired goodness score for a given future.

1.1.4 ADAPTIVE SEARCH OVER PROBE FITNESS LANDSCAPES

In another recent project (National Institute of Standards and Technology/Advanced Technology (NIST/ATP) Program-funded), we faced a problem similar to the PAINT challenge: Find a feasible set of changes (Course of Action / COA) to the current design of a car body (baseline polyagent simulation without any influencing actions) that improves its (simulated) assembly quality (outcome of Information Collection Plan / ICP). There we had available to us a detailed simulation of the assembly process of a car body with a set of parts and tools that may (active probe) or may not (passive probe) have been changed from the current prototype base design. We solved this problem in a swarming search process, where we designated a number of Search Agents (as many as we have simulation hosts available) that would explore the space of possible manipulations of parts and tools (COA’s) concurrently, trying out their “Solution Candidates” (current COA) against a complex simulation of the assembly process. Each Search Agent would try out a series of possible solutions (COA’s), evaluating their fitness based on a metric applied to the result of the simulated assembly process.




A Search Agent would hold a particular solution (COA), evaluate its fitness (relative to the baseline), compare it to previously evaluated COA's, and then decide how to change the solution (COA). Particular solution (COA) changes would either be short-range moves in solution space (minor changes to the plan to affect parts or tools) or long-range jumps (major changes to the COA). The moves would be determined by so called "Modifiers" that implement various "hill-climbing" heuristics, while the jumps offer qualitatively new solutions based on complex reasoning about the problem at hand performed by so called "Solvers". Search Agents are designed to prefer solutions from Modifiers until their series of changes did not result in a significant improvement against a previously attained solution for a while. Then the Search Agent would contact a Solver for a jump to a qualitatively new solution beyond the correlation distance of the search space. The probe generator discussed in the previous section would be such a Solver and we would have to adapt backtracking heuristics to develop at least one Modifier.


So, a Search Agent holding a particular COA with a given fitness would first decide whether to move or jump to its next COA. Then it would pick among the potentially many Solvers (agent went for a jump) or Modifiers (agent decided to slightly change its current COA) to select the one that would give it the next COA. Once it receives a new COA, it would run multiple instantiations of the Leadership/Policy system until it has sufficient input to determine the most appropriate ICP. Finally, it would assign a fitness value relative to the baseline ICP to its current COA and repeat the process.

1.2 POSSIBLE WORLD GENERATOR



In the following, we offer a series of annotated slides that show and motivate our draft for the architecture of the Possible World Generator, which is the low-dimensional stand-in for the integrated system of Component Predictive Models in the PAINT architecture.



The PAINT SEEDS system performs a resource-aware and adaptive search over the space of possible probes, evaluating the “fitness” (ability to reveal insights earlier) of each probe in an efficient multi-trajectory emulation of the PAINT Integrated Model.




To ensure a realistic fitness evaluation of each candidate probe in the multi-trajectory emulation in the Possible World Generator, we have to be careful in our construction of its abstract models. It is imperative that the emergent dynamics of the PAINT Integrated Model as seen through relevant observables are closely matched by those of the Possible World Generator in PAINT SEEDS. Without this correlation, there is no guarantee that the scoring of candidate probes in PAINT SEEDS is in any way correlated to the scoring of the same probes in the PAINT Integrated Model.





As a first step in ensuring the required correlation, we define the same high-level structure of models in the Possible World Generator as it is currently emerging in the program for the PAINT Integrated Model. We will have a leadership model, a pathway model, and, to ensure efficient execution, we limit the interactions between these two abstract and probabilistic models to sub-symbolic interactions.

The abstract and probabilistic models that are our stand-ins for the leadership and the pathway model in the PAINT Integrated Model are graph-based. These graphs provide topologies **within** which simplistic agents interact. The graph for the leadership model reflects the social network relations in the PAINT Integrated Model while the graph of the pathway model represents the technology advancements needed to achieve the leadership's goals.



1.2.1 THE PAINT SEEDS PATHWAY MODEL



Our abstraction of the pathway model represents technology advancement action (TAA) nodes that model attempts to advance particular technology development state (TDS) with a probabilistic outcome. Emulating an advancement action amounts to sampling two probability functions – one to determine the resulting state and one to specify the time it would have taken to reach this new state. The shape of these probability functions is determined by the amount of an abstract “resource” (e.g., funding, priority, staffing) that is currently allocated to this step. Pathway ghosts – probabilistic agents that are part of the NewVectors polyagent construct – transition through these nodes, sampling the probability functions.




The current resource level of all TAA nodes in the pathway model are set dynamically, relative to a changing estimate of what the current (“now”) state of the technology advancement is. Competing hypotheses of the most recently known technology development state are represented by Pathway avatars which issue pathway ghosts at that time and state at frequencies proportional to their respective certainty. The pathway ghosts then traverse the graph of TAA nodes, advancing in technology development state and time.




Each ghost deposits digital pheromones on the graph that are tagged with the ghost's current TDS and time stamp. Possible future technology development trajectories with a high likelihood receive many deposits while those of low likelihood are marked up only infrequently. Each ghost also deposits an untagged "pathway" pheromone used for normalization of pheromone concentrations into probabilities.

Once a sufficient number of ghosts have traversed the pathway model, we can identify the most likely technology development path and characterize it in regards to the expected technology development states and the time at which these states are reached.



1.2.2 THE PAINT SEEDS LEADERSHIP MODEL



We extract a social network graph from the (input to the) leadership model of the PAINT Integrated Model and instantiate a corresponding graph in the abstract and probabilistic leadership model of the PAINT SEEDS Possible World Generator. The nodes in this graph are entities, which are either individuals or organizations. The weighted links in the graph are their relations pertaining, in particular, to the possible flow of abstract “resources” among these entities.




The function of the social network in the PAINT SEEDS Possible World Generator is to emulate the flow of “resources” among the leadership. We postulate that there are some entities in the social network graph that have access to external resources (e.g., taxes, external companies). Furthermore, we assume that this influx of resources is time dependent.



As resources flow into the social network through distinguished nodes they are distributed in a process similar to spreading activation mechanisms in artificial neural networks. Thus, any node (indirectly) connected to a node that has access to external resources may receive a non-zero budget.

We associate an “entity avatar” agent with each node in the leadership model to determine the allocation of the available budget.



The avatar of an entity holds a target model developed by the SET (Science, Engineering and Technology) team in PAINT SEEDS. The target model represents a value judgment about future technology development states for the given entity, or, in other words, it describes the outcome of the technology development process (goal) that is desired by this entity. The role of the target model is to determine the allocation of resources available to the entity to technology advancement action (TAA) nodes in the pathway model. This decision is influenced by the pathway model's current estimate of likely technology development trajectories.

1.2.3 EXECUTING THE TWO MODELS

In the following, we discuss how the two models are executed together to create a multi-trajectory prediction of the technology development under the guidance by the possibly conflicting goals of the leadership.




FIGURE 17. STEPS 0-1.

Step 0: We initialize the abstract and probabilistic leadership and pathway models from the high-dimensional PAINT Integrated Model by digesting some of their input (e.g., structure of the alternative pathways), by observing their execution trajectory, and, potentially, by directly executing selected Component Predictive Models.

Step 1: We instantiate one or more pathway avatars at TDS/time locations representative of given hypotheses about the current or recent state of the technology development program under investigation. These avatars issue pathway ghosts that traverse the pathway model until their internal clock reaches the current time at which this analysis is performed. We consider the end-point of the most likely path of these ghosts the current state of the program, which defines the initial "now" of our abstract model.





FIGURE 18. STEPS 2-4.

Step 2: We execute the spreading activation process of resources in the social network, which determines the currently available budget of individual entities of the leadership.

Step 3: Using the target model in the entity avatars, we allocate resources to technology advancement actions in the pathway model, beginning at the current “now”.

Step 4: We emulate the evolution of the technology development program under the specified resource allocations by executing the pathway ghosts up to the final prediction horizon (e.g., 10 years into the actual future). The ghosts traverse the pathway graph, sampling the TAA nodes’ probability distributions, and depositing pheromones at the nodes.



Step 5: Analyzing the resulting probability distributions over the space of possible pathways, we determine the most likely pathway and characterize its timing and technology state sequence.

Step 6: Using the timing of the most likely pathway, we advance the temporal location of the model’s “now” marker one unit (e.g., a week) into the future.

Step 7: We communicate the probability distribution over pathways along with the most likely pathway to the leadership model.





FIGURE 20. STEPS 8-11.

Step 8: We repeat the spreading activation process of resources in the leadership model at the new location of the “now”, thereby updating the budgets available to the entities at the next decision cycle.

Step 9: The target models analyze high-probability technology development trajectories and their outcome and compare them with the goals of the entity they represent.

Step 10: Using the available budget, the entity avatars may attempt to change these trajectories by changing their allocation of resources to nodes in the pathway model. The entities are only allowed to change allocations to nodes that are still in the future relative to the current setting of “now”.

Step 11: We emulate the evolution of the technology development program under the specified resource allocations by executing the pathway ghosts up to the final prediction horizon (e.g., 10 years into the actual future). The ghosts traverse the pathway graph, sampling the TAA nodes’ probability distributions, and depositing pheromones at the nodes.



Step 12: Analyzing the resulting probability distributions over the space of possible pathways, we determine the most likely pathway and characterize its timing and technology state sequence.

Step 13: Using the timing of the most likely pathway, we advance the temporal location of the model's "now" marker one unit (e.g., a week) into the future.

Step 14: We communicate the probability distribution over pathways along with the most likely pathway to the leadership model.

We continue this repeated interaction of leadership and pathway model until the marker of "now" in the pathway model has passed the final prediction horizon. At this point, we have locked in all relevant resource allocations and thus the most likely trajectory through the pathway model would no longer be changed by the leadership model.


As the result of the execution of the two abstract models, the Possible World Generator returns a probability distribution over pathways that encode multiple more-or-less likely futures of the evolution of the technology development program under the guidance of the goal of the leadership. From this distribution we can now extract trajectory bundles in the Possible World Analyzer component of the PAINT SEEDS system.

1.3 ILLUSTRATION-OF-CONCEPT DEMONSTRATION

After defining the overall PAINT SEEDS architecture, we focused on implementing and demonstrating an illustration-of-concept prototype that shows the main processes of prediction and probe execution within the Possible World Generator component. Following the guidance by the program office, we chose a scenario based on nuclear capabilities development derived from open-source data guided by a generic set of actors.

1.3.1 PATHWAY MODEL

The nuclear technology demonstration pathway model was constructed with several goals in mind: several TAA Nodes in the pathway model graph have a non-zero starting set of resources committed to them; there are several TAA Nodes that are leaf nodes, meaning that they have no child nodes, and are therefore the end result of a particular technology path; there are several TAA Nodes that have multiple parents, meaning that they have many different paths they can play a part in.



There are many TAA leaf Nodes of note: Coal-fired Power Plant, Nuclear Power-plant, Depleted Uranium Munitions, Intercontinental Ballistic Missile, Medical Radiation Therapy, Nuclear Powered Satellite, and Nuclear Submarine. These leaf nodes were chosen for demonstration purposes because they represent the full range technology end-point possibilities: from the most benign (e.g. Medical Radiation Therapy) to technologies that represent a clear threat to the internal national community (e.g., Intercontinental Ballistic Missile).

The node pathways are highly interconnected, so that many TAA nodes play a key role in multiple technology pathways, representing varying degrees of threats. This was done to show how resources committed to one node or another could have dramatically different effects on the final pathway outcome (i.e., the Most Likely Path). Note also that because many TAA Nodes have multiple parents, they can be viewed as having multiple pre-requisites. For example, Intercontinental Ballistic Missile, the most threatening technology end-point, has four parents that are prerequisite TAA Nodes. In the current implementation of the pathway model World Generator, prerequisites are not considered for Most Likely Pathway determinations.

1.3.2 LEADERSHIP/TARGET MODEL

We have developed a user interface component that visualizes the leadership model and allows a user/analyst to make changes to the model. The leadership model is comprised of a set of generic leaders, such as President, Vice President, Minister of Defense, etc. The arrows between leaders represent the influence one leader has over another, which corresponds to how resources will be distributed among them.




FIGURE 23. THE NUCLEAR TECHNOLOGY DEVELOPMENT LEADERSHIP AND TARGET MODELS.

There are four pieces of information associated with each leader. The first is the leader's name, which appears below the leader's icon. The second is the set of bars to the left of the leader. These represent the resources a leader has, and each bar corresponds to a type of resource, such as money, political power, military resources, etc. Although there are four types visualized, the system is currently only working with one type of resource. To the right of the leader is another set of bars, and these represent the leader's desires or goals. Each leader has the same set of bars, but the levels of the bars indicate how interested the leader is in each goal. In this setup, each bar corresponds to an end node in the pathway model. We have chosen 4 goals a leader may be interested in: two benign goals (medical radiation therapy and electromagnetic isotope separation) and two nefarious goals (intercontinental ballistic missile and nuclear submarine). The user can adjust a leader's goals by dragging the bars up or down. In this example, most leaders are primarily interested in the benign goals, but the enemy leaders are primarily interested in the nefarious goals. The fourth piece of information associated with a leader is his satisfaction level, which appears as a percentage above the leader icon. This value indicates how well the most likely path in the pathway model matches up with the leader's desired goals.

The interface sports additional capabilities to adjust the leadership model, but currently have no effect when the system starts running. This includes the ability to add/remove leaders and to edit the links between them. The user/analyst can also adjust the resource bars to the left of the leader, but for now, the resources are controlled by the pathway client, which does not grab resource values from this interface.

The user hits the Publish button to send preference information to the pathway model. The preference information contains the resource allocation strategy for each leader. Each leader is to put a percentage of his resources on a select set of nodes in the pathway model. The allocation strategy is determined by each leader's goal preferences. A leader will only put resources on the goal he cares about most. The percentage of resource he parts with is proportional to how much he cares about that top goal. Of the resources he devotes towards the goal, a fraction of it will be put on each node in the pathway model that is on a path that leads towards the goal.

After the preference information is published, the pathway client will compute the most likely path and return the most likely path information to the leadership model. When the leadership model receives this message, it will update the satisfaction levels of the leaders. If the most likely path leads to a leader's top priority goal then that leader's satisfaction level will be 100%. A leader who did not have his top goal met, but still cares about the goal in the most likely path, will have a satisfaction level proportional to his interest in the goal. A leader who had no interest in that goal will have a 0% satisfaction level.

2 CONCLUSION

Within the available funding for the PAINT SEEDS FY'07/'08 work package we accomplished two main goals. We defined a detailed and realistic architecture for a generate-and-test approach to probe recommendation and we illustrated the main concepts of the execution of probes in a low-dimensional multi-future polyagent simulation and the prediction of the outcome of the interaction of the leadership goals with the constraints of technology development. We now stand ready to take the project to the next level towards a full implementation of all the components and processes laid out for the full system.

3 LIST OF ACRONYMS

PAINT – Pro Active Intelligence

SEEDS – Simulated Exploration of Executable Design Strategies

CPM - Component Predictive Models

BAE – British Aerospace

LCC- Language Computer Corporation

NLP – Natural Language Processing

NIST- National Institute of Standards and Technology

ATP – Advanced Technology Program

COA – Course of Action

ICP – Information Collection Plan

TDS – Technology Development State

TAA – Technology Advancement Action

SET – Science, Engineering and Technology

PWG - Possible World Generator