REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*: 08-12-2006

2. REPORT TYPE: Final

3. DATES COVERED *(From - To)*: 11-01-2002 to 12-31-2005

4. TITLE AND SUBTITLE

Adaptive Architectures for Effects based Operations

5a. CONTRACT NUMBER: N00014-03-1-0033

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

6. AUTHOR(S)

Alexander H. Levis

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

System Architectures Laboratory
George Mason University
ECE Dept., MS 1G5
Fairfax, VA 22030

8. PERFORMING ORGANIZATION REPORT NUMBER

SAL-06-501-R

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Office of Naval Research
Code 342
825 N. Randolph St.
Arlington, VA 22203-1995

10. SPONSOR/MONITOR'S ACRONYM(S)

ONR

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT

Approved for Public Release; Distribution is Unlimited

13. SUPPLEMENTARY NOTES

----

14. ABSTRACT

The report documents the research on adaptive architectures for effects based operations. The objectives were to provide a predictive capability against transnational threats and to expand the capability of tools for developing courses of action that include both kinetic and non kinetic actions. The theory and algorithms of Influence nets were developed further to include timed influence nets and dynamic influence nets. The research results were implemented in a tool called CAESAR II/EB (now called Pythia) and were used in several simulated and real applications. The applicability of the approach to transnational threats such as terrorist organizations was investigated. An extension of previous work on organization design was made to incorporate the notion of task graphs.

15. SUBJECT TERMS

Influence nets, Timed Influence nets, Effects based operations, Organizational Architectures

16. SECURITY CLASSIFICATION OF: a. REPORT	b. ABSTRACT	c. THIS PAGE	17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
Unclas	Unclas	Unclas	None	117	Alexander H. Levis

19b. TELEPHONE NUMBER *(Include area code)*: 703 993 1619

Standard Form 298 (Rev. 8/98)
Prescribed by ANSI Std. Z39.18

GEORGE MASON UNIVERSITY
SCHOOL OF INFORMATION TECHNOLOGY AND ENGINEERING
DEPARTMENT OF ELECTRICAL ENGINEERING
SYSTEM ARCHITECTURES LABORATORY

ADAPTIVE ARCHITECTURES FOR EFFECTS BASED OPERATIONS

FINAL TECHNICAL REPORT

Contract No. N00014-03-1-0033

for the period

1 November 2002 to 31 December 2005

20060818015

Submitted to:
Office of Naval Research
Attn: Mr. Gerald Malecki
Code 342
825 N. Randolph Street, Suite 1425
Arlington, VA 22203-1995

Prepared by:

Alexander H. Levis
Principal Investigator

August 12, 2006

TABLE OF CONTENTS

SECTION 1

Executive Summary

1.1 Introduction

This report documents the research conducted by the Systems Architecture Laboratory, Dept. of Electrical and Computer Engineering, George Mason University (GMU) on Adaptive Architectures for Effects based Operations. While the project was based on work done under the previous Adaptive Architectures for Heterogeneous Command and Control (A2HC2) grant, it contained new research directions that reflected the effort to transition some of the research work to the operational Navy. Indeed, a new task was added late in the program in support of Operation Iraqi Freedom.

There are many new challenges facing the country, the Department of Defense, and the Navy in the post September 11, 2001 era. They include the ongoing challenges of organization, breadth and frequency of military operations, and joint, combined, and coalition operations. We still face the problems of designing and understanding the behavior and performance of collaborating command centers, particularly in the context of network centric operations. We need to conduct effects based operations, and, in addition to traditional operations, we face an adversary that conducts asymmetric operations with an untraditional organization and structure that is transnational and that is difficult to detect, monitor, and, if necessary, defeat. While these transnational threats have been in existence for some time, the events of 9/11 have elevated their importance to the highest priority.

One of the newer constructs that the Navy has introduced is that of the Expeditionary Strike Group (ESG). It is transformational in that it focuses on creating a leaner, more mobile and flexible force that is capable of responding to a wide spectrum of conflicts and contingencies. Missions of the ESG range from humanitarian assistance and disaster relief, to power projection ashore via amphibious or airborne Marine assaults or cruise missile strikes, to Maritime interdiction operations. The ESG is also capable of integrating into or leading a joint operation involving some or all of the other services of the United States Armed Forces.

Previous research, sponsored by the Office of Naval Research, has provided a foundation for addressing the dynamics of adaptation in a network centric environment we are observing in war games and experiments as well as in the operational world. These phenomena include the adaptation of the organizational structure and procedures of the entities that make up the command and control system to take advantage of the expanding capabilities afforded by emerging information technology tools.

The naval services are applying the principles of effects based operations instead of the more traditional attrition based approaches. For example, the battle plan in Global 2001 was based on Effects Directives that were directly linked to commander's intent and supported by models and analysis using the CAESAR II/EB Course of Action (COA) analysis tool (now called Pythia) and techniques that were developed further in this project. The effects based approaches require a good understanding of the adversary[1]. The problem of understanding and modeling the adversary has become more challenging as a result of the transnational threats that have emerged.

Because of the importance and challenges presented by transnational threats, we proposed a new thrust in our research activities. We needed to develop tools and techniques that provide the capability to anticipate the potential actions of a transnational adversary so that appropriate actions can be taken. The need for this predictive capability adds complexity to the Effects Based Operations problem that we have been addressing using influence nets and discrete event system models. In our preliminary assessment and formulation of the problem, we postulate that transnational threats can be thought of as having both an operational architecture and a non-traditional systems architecture. By using our modeling and analysis tools and techniques to develop and understand the operational architecture, it may be possible to infer portions of the systems architecture. The systems architecture can then be exploited to obtain information about the status of operations within the transnational organization. Thus, an important new thrust was to formalize a procedure for modeling transnational threats as loosely-coupled, network-type organizations so that effects based analysis can be preformed.

We have learned a great deal about using the tool we developed, named Pythia, to support the planning of effects based operations. The lessons learned have suggested what new capabilities should be developed. In particular, we recognized the need to extend the tool capabilities and develop the procedures to support not only effects based Course of Action development, but also effects based execution and assessment.

1.2 Objectives

We proposed four main thrusts(objectives). The first objective was to develop the theory and modeling techniques to provide a predictive capability against transnational threats. The second objective was to continue to expand the capability of the effects based operations tools to support the integrated Course of Action planning of lethal, non-lethal (e.g. information operations), and ISR operations, and monitoring and assessment of operations using feedback from resources and ISR sensors. The third objective was to continue the collaborative efforts with the Naval Post Graduate School, the University of Connecticut, Carnegie Mellon University, and others in the model driven experimentation paradigm that has been employed successfully in A2C2. Finally the fourth objective was outreach research, which evolved in supporting the Expeditionary Strike Group through a series of interactions.

[1] Wagenhals, L.W., "An Operational Concept For Effects Based Course-Of-Action Development And Evaluation In A War Game", Proc. of the 2002 SPIE Aerosense Conference, Orlando, FL, April 2002

1.3 Tasks

The scope of work of this research project consisted of four tasks; a brief summary of the work done in each task is included.

Task 1: Effects Based Operations against Transnational Threats

A terrorist network can be described in terms of its operational and system architectures but the mapping between these architectures is less well known and understood since the operational architecture can be mapped into numerous system architectures that are flexible and reconfigurable and contain target sets that are both hard and soft targets such as political, religious, social and economic networks. Traditional attrition-based warfare focuses on destroying the hard targets of the system architecture of the adversary but terrorists are very much unlike the military forces modeled in force-on-force type engagements and hence, to suppress, if not destroy, transnational terrorism it will be necessary to attack and destroy not their system architecture but their operational architecture–their ability to conduct operational activities in support of their goals. The concept of effects based operations lends itself well to modeling and assessing approaches to destroying, degrading or disrupting terrorist acts. The CAESAR II/EB research tool has been used to construct influence nets and courses of action to mitigate terrorist attacks. The findings from this exploratory research are presented in Section 2 of this report. This is work in progress and much remains to be done.

Task 2: Models to support effects based operations

At the core of the effects based modeling approach is the modeling of courses of action to produce the desired effects by using Influence nets. In earlier work, the static influence net theory was extended by introducing Timed Influence nets (TINs). Section 3 presents an algorithm for transforming Timed Influence Nets into TimeSliced Bayesian Networks (TSBN). The advantage of TINs lies in their ability to represent both causal and time-sensitive information in a compact and integrated manner. They are used to help a decision maker model the causal and temporal interdependencies among variables in a system. The TIN formalism offers a suite of analysis tools that can be used by a user to analyze the impact of alternate courses of actions on likely outcomes. An even larger, and more robust suite of analysis tools exists for TSBNs. These algorithms also allow analyses that are not available in the TIN formalism, e.g., provision for incorporating real-time information in the form of evidence regarding certain variables and calculating its impact on the rest of the system. The knowledge acquisition process of TSBNs, however, is intractable for large models. This work constitutes an attempt to combine the advantages of both modeling paradigms, TIN and TSBN, into a single formalism by providing a mapping from a TIN to a TSBN. The new formalism uses the TIN approach for the model building and the TSBN for analysis and evaluation. A system analyst, in this combined approach, interacts with a TIN, and the analysis results obtained on the TSBN are mapped back to the TIN, making the transformation completely hidden to the analyst.

In a second sub-task, an approach for belief updating in Timed Influence Nets was developed. Influence Nets provide graphical representation of causal or influencing relationships in complex situations. They are used to model and evaluate courses of actions in certain domains and to compare the performance of actions based on the desired outcome. In Timed Influence Nets, the impact or effect of these actions on target variables is not instantaneous. This is modeled by adding communication and processing delays in the model. The approach provides a technique for updating the beliefs of variables in the model over time once new evidence is received about some of the variables in the model. The objective is to assess the behavior of the variables of interest as a function of both the timing of actions and the receipt of evidence on indicators, thus providing aid to decision makers in the revision of the planned courses of actions. The results of this effort are described in Section 4.

A further research effort focused on developing structural and parametric enhancements to the Timed Influence Nets. The existing TIN framework does not have the capability to model the impact of different sequences of actions. Thus, no matter what the sequence of action is, the final outcome remains the same. Furthermore, it is assumed that the influence of an event on another event is stationary, i.e., the influence remains the same throughout the campaign. Both of these constraints may turn out to be unrealistic in many real world situations. The enhancements introduced in Section 5 of this report would overcome the above two limitations. The proposed structural enhancement would enable a system modeler to model the impacts of different sequences of actions on the desired effect; while the parametric enhancements would aid the mathematical modeling of time-varying influences. Together these enhancements make it possible to model the impact of repetitive actions in a dynamic uncertain situation.

In Section 6, an Evolutionary Algorithm (EA) based approach for finding effective courses of action (COAs) in a complex uncertain situation is described. The complex situation is modeled using the TIN probabilistic modeling and reasoning framework. The TIN-based framework helps a system modeler in connecting a set of actionable events and a set of desired effects through chains of cause and effect relationships. Once a TIN is built, the optimization task confronted by the modeler is to identify a course of action that would increase the likelihood of achieving the desired effects over a pre-specified time interval. This research task used EAs to accomplish the optimization. The proposed approach generates multiple COAs that are close enough in terms of achieving the desired effect. The purpose of generating multiple COAs is to give several alternatives to a decision maker. Moreover, the alternate COAs could be generalized based on the relationships that exist among the actions and their timings of executions. While determining an effective course of action in a given situation, a system modeler has to consider several temporal/causal constraints that are present in a problem domain. The approach also includes a constraint specification language that would help a system modeler in specifying these constraints.

Task 3: Model driven experimentation

Future command and control centers are being designed to exploit the capabilities offered by information technologies; models of these proposed architectures are necessary to predict the performance of alternative designs. However, many of the performance metrics that these future centers will be evaluated on, such as speed of command and shared situational awareness, have not been included in previous command and control models. An enhanced command center model has been created that combines both a task process model and a decision maker model in order to provide the necessary degrees of freedom required to evaluate such performance metrics. The task process model was developed as part of a recent pre-experimental modeling iteration for a subject experiment and captures the stages of a task over its lifetime; the decision maker model is required in order to explicitly represent the interaction between decision makers as they process the task. This enhanced model will allow sophisticated modeling of interactions between decision makers, such as decision maker synchronization and information sharing. By combining the task process model with the decision maker model, surrogate measures of speed of command and situation awareness can be developed and used to evaluate the behavior and performance of command and control information and decision processes, essential to assess any future command and control architecture. The results of this task are described in Section 7.

Task 4: Outreach Research

The outreach effort had two components. The first one was focused on providing support, as part of the A2C2 Program Team, to the Expeditionary Strike Group One as it prepared for deployment in the summer of 2005. This included visits with the ESG 1 Commander, Rear Admiral Michael LeFever, and his senior staff and presentation of tutorial lectures on Effects based operations. The second effort was carried out at the invitation of the Joint Improvised Explosive Devices Defeat Organization (JIEDDO). We were asked to use the Pythia tool, together with a data set provided by DoD, to explore the tools use for developing courses of action that include both kinetic and non kinetic operations that take place over a long period of time. This effort started at the end of the reporting period covered by this report and continued under other sponsorship after that. The results of this ongoing task have not been documented yet.

1.4 Conclusion

While substantial progress has been achieved in all tasks, the problem of developing adaptive architectures that will enable the warfighters to conduct effects based operations against non-traditional adversaries remains challenging. Some of the challenges have to do with the nature of the threat, some with the environment, and some with the resources that need to be applied. The research conducted under this contract has established some foundations for further research and has led to the development of tools that address specific aspects of the problem.

SECTION 2

Effects Based Operations for Transnational Terrorist Organizations: Assessing Alternative Courses of Action to Mitigate Terrorist Threats

Larry K. Wentz and Lee W. Wagenhals

2.1 Introduction

Transnational terrorism is a multidimensional problem for which motivation is a key enabler. Terrorists are inspired by many different motives, some rational but most not, and they have goals. Some terrorists are rational thinkers and they carefully assess whether they can induce enough anxiety to attain their goal without causing a backlash that will destroy the cause and the terrorist themselves. Others may be motivated for psychological reasons that are derived from personal dissatisfaction with their life or accomplishments. Culture is another key motivator and in this regard, there is a tendency for western societies to reject, as unbelievable, things such as vendettas, martyrdom and self-destructive group behavior. Terrorism thrives in a sea of perceived injustice and religion is probably the most volatile of culture identifiers.

Security is another important consideration that influences terrorist organizational arrangements (cellular structures seem to dominate) and recruitment and training (tend to be extremely security-sensitive activities). There is a strong incentive by the members of the networks to keep their structure and operations secret and unobservable. As a result, intelligence operations against these organizations and their leaders, members and supporters are extremely complex and difficult. Terrorist communications are multidimensional and include means such as email, Internet web sites, commercial telecommunications, cellular, courier, radio/TV and other covert or non-traditional means. They use the mass media to generate fear and panic in a free-minded public and also exploit the global media and information highways to carry news of their violence along with propaganda of the deeds. On the other hand, media coverage of terrorism by the free world can be used to educate the public, temper public anxiety, and influence actions to prevent and counter terrorist actions.

Transnational terrorist networks are hard to define in terms of geographical boundaries or through their physical assets. What characterizes these networks is not so much their system architecture but their operational architecture. Inactive nodes can come to life temporarily to carry out an operation at some location and then may go inactive again or self-destruct. Or, in some cases, a system node may augment itself with additional physical assets to carry out an operation and then discard these assets or disengage from them. At the operational level, the relationships that tie the network together, the interconnections, can be a set of beliefs, a financial infrastructure and a communications infrastructure. It is, therefore, dangerous to see them only as madmen bent on destruction.

The terrorist are very much unlike the military forces modeled in force-on-force type engagements where traditional attrition-based warfare focuses on destroying the system architecture of the adversary and the relationship (mapping) between the operational and system architecture is well known and well understood.

The terrorists deliberately avoid engaging enemy military forces in combat and do not function in the open as armed units. For the terrorist network, the operational architecture maps into numerous system architectures. Therefore, an important objective in suppressing, if not destroying, transnational terrorism is to attack and destroy not their system architecture but their operational architecture—the ability to conduct operational activities in support of their goals.

For military opponents, a well defined mapping between the operational and system architecture leads directly to concepts such as physical Centers of Gravity, prioritized target lists and the like. But, when the adversary is characterized primarily by an operational architecture that maps into many system architectures or to flexible system architectures that can be easily reconfigured, there is a need to change the way they are analyzed and modeled. The concept of effects based operations is well suited to addressing this problem. Instead of focusing on the servicing of a well-defined a priori target list, the focus is on the effects to be achieved. The target list still exists and includes both hard and soft targets: from weapons systems, to C2 nodes, to leadership nodes, to infrastructure nodes, to political, social, and economic nodes, to the contents of communications, information, and databases. But, the target list is only an intermediate construct, a means to an end that can change rapidly as effects on the adversary are achieved or not. Indeed, the list of possible actions to be used against the adversary centers of gravity (political, military, economic, social, information, and infrastructure) includes all instruments of national (or coalition) power: diplomatic, information, military, and economic. The availability of all instruments gives added flexibility in trying to achieve the desired effects and to avoid undesirable ones. But, it also makes the Course of Action (COA) problem and the subsequent planning problem much harder. There are now many alternatives, many choices. The choice of a set of actions, their sequencing, and their time phasing become problems in their own right.

Hence, effects based operations for transnational terrorism threat mitigation requires not only a deep understanding of the terrorist motivation, methods, organization and other factors but also needs an understanding of the friendly capabilities and infrastructure and likely vulnerabilities that might be of interest to terrorist. Additional work needs to be done to develop a more informed understanding of the appropriate relationships of motivators, organization dynamics and capabilities of terrorists and courses of action. There are a number of tools that address pieces of the problem but the current suite of tools available in the community does not fully address an integrated approach to counter terrorism course of action planning and assessment.

During the George Mason University (GMU) support to the Joint Forces Command sponsored Millennium Challenge 2002 experiment,[2] an attempt was made to use the GMU effects-based course of action planning and assessment research tool, called CAESAR II/EB, to construct an influence net for developing and assessing courses of action to deter a terrorist attack within the region of blue force operation for the experiment. The results of this effort were used in support of follow-on GMU research into developing influence networks to examine courses of action that might be considered to deter an act of terrorism. Findings from literature searches and other research activities have been used as an integral part of the research effort presented herein. Documents on the *Terrorism Research Center* Internet web site (www.terrorism.com) and RAND publications by Bruce Hoffman and Brian Jenkins were particularly helpful as were the numerous other documents listed in the References. These information sources were used extensively to develop the terrorism insights needed to build the case study model presented herein. Based on principles set forth in the US "*National Strategy for Combating Terrorism*," alternative high-level courses of action that brought to bear elements of national power were developed and assessed using the case study model. This paper explores some of the challenges of developing and assessing EBO courses of action to mitigate terrorist threats and provides an example of a counter terrorism influence net and some findings from an assessment of COAs aimed to prevent terrorist actions. This is work in progress and much remains to be done.

2.2 CAESAR II/EB, The Tool

The CAESAR II/EB tool was originally designed to support the analysis of an adversary's actions and reactions to Blue's activities so that COA options could be evaluated in a rigorous manner. It was inspired by the need to support the development of Information Operations (IO) influence planning and its integration with traditional military operations. The tool incorporates influence nets as a probabilistic modeling technique and a discrete event system modeling technique, Colored Petri Nets (CP net), to support the temporal aspects of COA evaluation. These two techniques enable the modeler to create the structure of actions, effects, beliefs and decisions and the influencing relationships between them. The strength of the influencing relationships is also captured. The influence net provides a static equilibrium probabilistic model that indicates the probability of effects given sets of actions. A mapping has been established and an algorithm has been encoded for automatically converting the influence net to a CP net. After an influence net is converted to a CP Net, temporal analysis can be conducted that provides the probability of effects over time given a timed sequence of actions. This tool was designed to develop and assess COAs at the operational and strategic level.

The influence net provides an environment for modeling of the causal and influencing relationships between actions by our forces (Blue) and effects on the adversary (Red). It uses a graphical representation comprised of nodes that represent actions or effects and causal or influencing relationships between the nodes. In addition to the network structure of the model, estimates of the "strength" of the causal and influencing

[2] Part of this work was supported by the Air Force Office of Scientific Research under grant No. F49620-02-1- 0332

relationships is added and enables an underlying probabilistic model base on Bayesian mathematics to be used for analysis. The construct shown in Figure 1 is used. Starting from the set of desired and undesirable effects that reflect the goals of the mission, analysts work backwards to relate the effects to actions that are under our control. Once the Influence net has been completed, it can be used to evaluate the impact of actions on the effects (decisions) of interest using its underlying Bayesian mathematics.

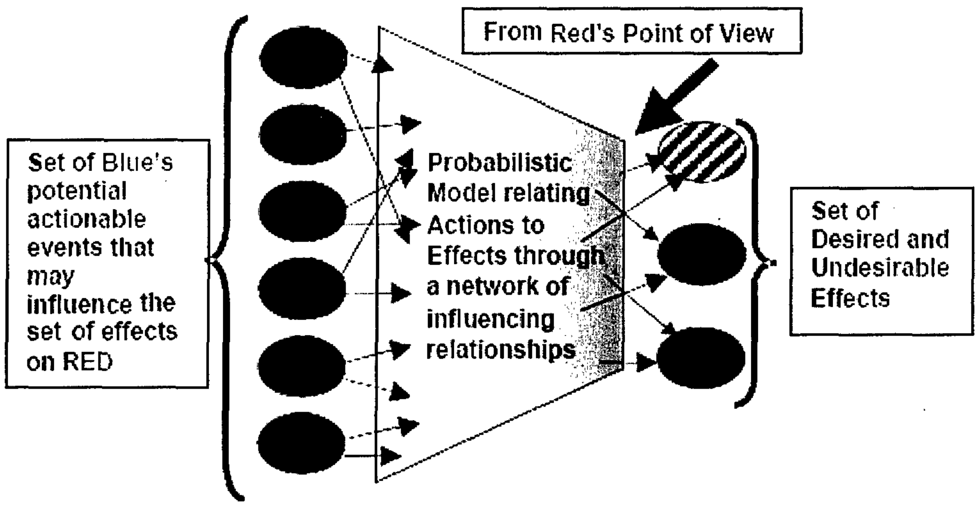

Figure 1. Modeling Actions and Effects

Once the analysis of the Influence net has been completed and the actionable events for the COA have been selected, planners assess the availability of resources to carry out the tasks that will result in the occurrence of the actionable events. The resultant plan will indicate when each actionable event will occur. Clearly, it is not only the selection of the set of actions that will lead to achieving the overall desired effects while not causing the undesired ones that is important. The timing of those actions is critical to achieving the desired outcomes. An algorithm has been implemented[3] that converts an influence net into a discrete event dynamical system model. The particular mathematical model used is that of CP Nets and their software implementation in Design/CPN 5[4]. The nodes in the Influence net become transitions in the CP Net and the places hold tokens that carry the marginal probabilities. Since the Influence net does not contain temporal information, it must be provided as an input to the CP Net.

Figure 2 shows the combination of models and results produced by the CAESAR II/EB tool. An Influence net model for a given situation is shown in the upper left of Figure 2. Each node represents an action, event, belief, or decision. A declarative sentence in the form of a proposition is used to express the meaning of each node. The directed arcs

[3] Wagenhals, L. W., Shin, I., and Levis, A. H. (1998). "Creating Executable Models of Influence Nets with Coloured Petri Nets," *Int. J. STTT*, Springer-Verlag, Vol. 1998, No. 2, pp. 168-181.

[4] Jensen K. (1997). *Coloured Petri Nets: Basic Concepts, Analysis Methods and Practical Use. Volumes 1, 2, and 3.* Monographs in Theoretical Computer Science, Springer-Verlag, Berlin, Germany.

between two nodes mean that there is an influencing or causal relation between those nodes. The truth or falsity of the parent node can affect the truth or falsity of the child node. The Influence net has been arranged with potential Blue actions on the left and the key Red decisions on the right. This is to indicate visually that the effects of the actions are expected to propagate to intermediate effects over time until their impact reaches the key decisions. This captures the cascading and accumulation of effects. There are six actionable events on the left side of the Influence net. These are candidate actions (or results of actions) that can comprise a COA that can impact the three Red decisions of interest.

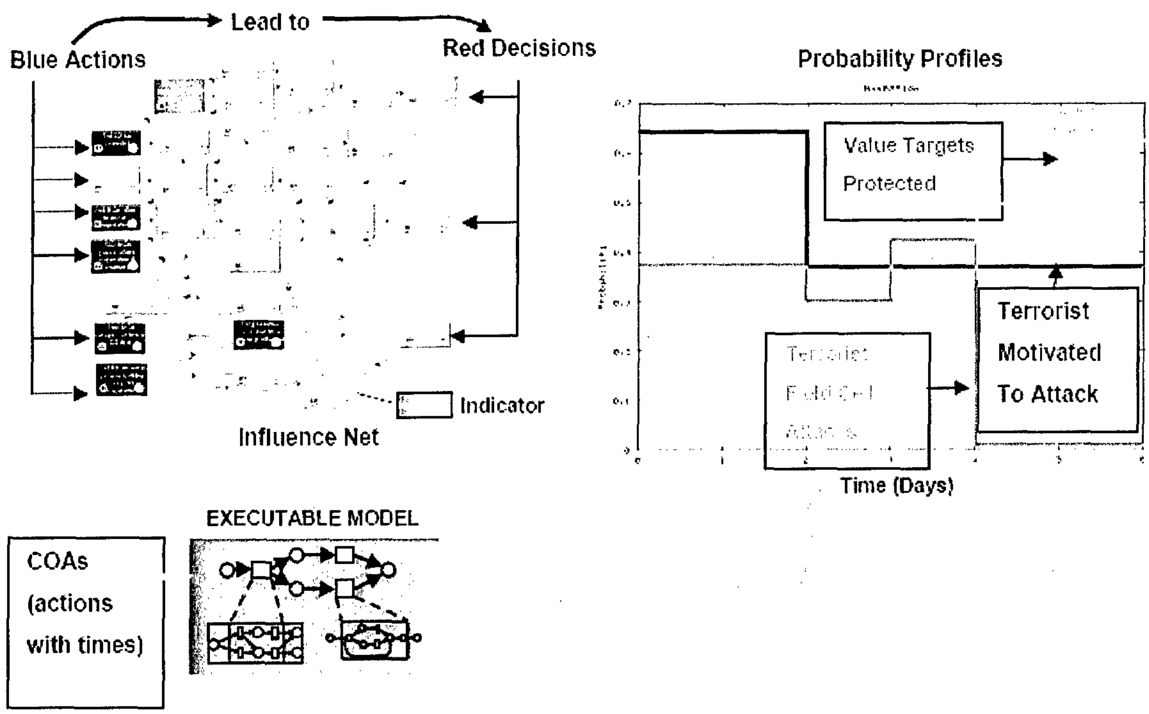

Figure 2. CAEAR II/EB Products

Once the analysis of the Influence net has been completed and the actionable events for the COA have been selected, the Influence net is automatically converted to an executable model (CP net) so that a temporal analysis of the COA can be performed. Using the executable model, the analyst is able to generate the probability profiles that show the marginal probability for any node in the net as a function of time. These profiles can indicate how long it will take for the effects of the actionable events to affect various nodes in the Influence net. The analyst will most likely concentrate on the probability profiles of the key decision nodes, the nodes with no children. The probability profiles shown in Figure 2 were generated for the COA proposed by the planners. The annotations have been added to indicate the three separate probability profiles. Different timing of the actions can alter the probability profiles. As a result, some will be more desirable than others while others may be unacceptable, so the planners will try to adjust the scheduling of actions.

2.3 Terrorism Definitions

There are numerous definitions for terrorism. The *U.S. National Security Strategy* defines terrorism as simply “premeditated, politically motivated violence against innocents.” U.S. government organizations and the UN define terrorism slightly differently.6[5] For example:

> **U.S. Department of Defense**: The calculated use of violence or the threat of violence to inculcate fear; intended to coerce or to intimidate governments or societies in the pursuit of goals that are generally political, religious, or ideological.
> **U.S. Department of State**: Premeditated, politically motivated violence perpetrated against noncombatant targets by sub-national groups or clandestine agents, usually intended to influence an audience.
> **U.S. Federal Bureau of Investigation**: The unlawful use of force or violence against persons or property to intimidate or coerce a government, the civilian population, or any segment thereof, in furtherance of political or social objectives.
> **United Nations**: A unique form of crime. Terrorist acts often contain elements of warfare, politics and propaganda. For security reasons and due to the lack of popular support, terrorist organizations are usually small, making detection and infiltration difficult. Although the goals of terrorism are sometimes shared by wider constituencies, their methods are generally abhorred.

The challenge of grasping the nature and parameters of the war on terrorism is certainly not eased by the absence of a commonly accepted definition or by its depiction as a Manichaean struggle between good and evil, “us” versus “them.”[6] Consensus on the definition of terrorism is not necessary to conduct counter terrorism operations against specific terrorist organizations but a lack of consensus can impede the study of the phenomenon itself.

Counter terrorism is not war in the traditional sense of military operations between states or between a state and an insurgent enemy for ultimate control of that state. Terrorist organizations do not field military forces as such and are trans-state organizations that are pursuing non-territorial ends. As such, and given their secretive, cellular, dispersed, and decentralized “order of battle,” they are not subject to conventional military destruction.

Based on findings from the research of literature on terrorism, what terrorism is and is not can be summarized as follows:

- **Terrorism is:**

– Calculated use of covert criminal violence or threat of violence
– Deliberately selected as a tactic to effect change
– Targeting of innocent people, including military personnel
– The use of symbolic acts to attract media and reach a large audience

[5] “The Terrorist Recognition Handbook” by Malcolm Nance (2003)
[6] Record, Jeffrey (2003). “Bounding the Global War on Terrorism,” Army War College, Strategic Studies Institute.

- Illegitimate combat, even in war
- Never justified

- **Terrorism is not:**

- Common crimes
- Conducting acts legal under national and international law
- Civil disturbances or spontaneous rioting
- Freedom of speech or nonviolent civil disobedience
- Protests and assembly to present opposing views and express dissent

The following terms are used by U.S. organizations such as the Defense Department, Intelligence Agencies and the Law Enforcement community to describe classes of measures taken to address terrorist acts.

Antiterrorism: Defensive and preventive measures taken to reduce vulnerability to terrorist attacks.

Counter-terrorism: Offensive measures taken in response to a terrorist attack, after it occurs.

Combating terrorism: The U.S. government program against terrorism includes antiterrorism, counter-terrorism, and all other aspects of tracking, defense, and response to terrorism throughout the threat spectrum.

Force Protection: The U.S. DOD program for the defense of military and government assets from terrorist and unconventional warfare attack—detect, deter, and defend.

2.4 Terrorist Considerations

Numerous reports from the *Terrorism Research Center* Internet web site and books and articles published on the subject of terrorism were used to develop the insights presented herein. Of particular value were the following:

- Terrorism Research Center Internet web site
    - The Basics: Combating Terrorism, an essay from the U.S. Army Field Manual 100-20, Stability and Support Operations
    - Terrorist Intelligence Operations, reprint from the Interagency OPSEC Support Staff, Intelligence Threat handbook

- Microsoft Encarta Online Encyclopedia 2003
    - Terrorism by Bruce Hoffman

- RAND
    - "Countering al Qaeda" by Brian Jenkins
    - "Countering the New Terrorism" by Ian Lesser and et al

- "Deterrence and Influence in Counter Terrorism" by Paul Davis and Brian Jenkins

– Books
- "Inside al Qaeda" by Rohan Gunaratna
- "Inside Terrorism" by Bruce Hoffman
- "Terrorism, War and the Press" by Nancy Palmer
- "The Terrorist Recognition Handbook" by Malcolm Nance
- "Framing Terrorism" by Norris Pippa and et al

Terrorists prefer simple strategies that appear sophisticated but are simple in planning and execution. They seek dramatic and wide publication by media to transmit fear and publicize their cause. Their apparent lack of logic enhances the terror in terrorism. Terrorist acts are seemingly random and they feel their goal will be reached by conducting enough attacks. They achieve their most dramatic impact through the use of speed, surprise and violence of attack. The terrorist only needs to get lucky once but the antiterrorist forces need to be lucky all of the time.

The goals of the terrorist organizations focus on recognition, coercion, extortion, intimidation, provocation, and insurgency support for their cause. Their objectives are to create a climate of fear in a targeted group or nation through a sustained campaign of violence and to destroy the social and political order by attacking and destroying commerce, property and infrastructure. They seek revenge for previous incidents or situations affecting terrorist organizations or its causes and try to negatively affect processes that the terrorist organization sees as against its interests. Attempts are made to eliminate specific individuals or groups and to demonstrate the weakness of legitimate governments. Terrorist organizations try to ensure governments overreact and oppress their own people. They continuously try to gain new recruits, money or weapons. Some terrorist organizations attack just to achieve the satisfaction of harming their enemy. Attacks also serve to demonstrate that the terrorist group is still active.

Terrorist groups can be indigenous or transnational. They can be state-sponsored, state-directed or have no state relationship. Those organizations that are state-sponsored tend to operate independently but receive support such as weapons, training, money, and safe havens. Those that are state-directed, act as agents of the state and receive intelligence, logistics and operational support. The groups not sponsored act autonomously and receive no significant support.

Motivation is a major consideration in terrorist organizations. Some are rational and think through goals and objectives, conduct course of action planning and assessments and risk and cost benefit analysis. They are careful when inducing anxiety to achieve their goals to attempt to ensure that it does not cause a backlash that may destroy them or their cause. Others are psychologically motivated and are dissatisfied with life and accomplishments and crave violence to relieve anger. They tend to need to belong to a group and require group acceptance, demand unanimity, are intolerant of dissent, and have a polarized "we versus them" outlook. Culture is another key motivator. Western cultures are reluctant to appreciate the intense effect of culture on behavior. In their view irrational behavior as a means to achieve objectives is counter culture. They believe rational behavior guides human actions and reject the notions of vendettas, martyrdom, self-destructive group behavior, and dissolution of a viable state for ethnic purity. For the terrorists, fear of cultural extermination leads to violence — the perception that "outsiders" are against them. Religion can be the most volatile of cultural identifiers — the belief in moral certainty and divine sanctions.

Security is a primary concern of terrorist organizations. Although cell operations are the least understood part of terrorism, it is believe terrorist organizations are best served by cellular structures that operate in secret as small team. This way, members do not know and cannot identify more than a few the other members. They can operate as a group on orders of a commander or independently. Defections are rare and it's difficult to penetrate cells. Fundamental units such as Command and Control, Tactical Operations, Intelligence, and Logistics are employed. A highly trusted and experienced leader generally runs the Intelligence cell and members of this cell rarely participate in attacks — there is a need to protect identity of members. Terrorists tend to organize to function in the environment where they plan to carry out their attacks — this is situation specific. Numerous means are used to communicate. Direct means such as face-to-face, Internet, cell phones and telephones can be used. Indirect means such as courier, trusted agent, Internet, cell phones, telephones, mail, dead drops, newspapers, books, and Radio/Television are used as well. Charismatic leaders are needed to unite the effort otherwise behavior is a reflection of the group dynamics. The support structure is a mix of state-sponsors and sympathizers. The recruitment process is highly security-sensitive. Training of the terrorist organizations can vary from military style at sophisticated facilities to inspirational talks before activation — motivating "throw away" operatives.

Terrorist potential targets generally fall into hard targets that are security conscious and difficult to attack successfully and soft targets that are people, structures, or locations that have less security and are open to public. Target selection is based on motive (ultimate goal/objective), opportunity (feasibility) and means (covert capabilities). The targets they choose can be categorized as follows:[7]

- **Strategic value:** Long-term impact target sets that include executive leadership, strategic reserves, cities, and national command centers.
- **High payoff:** Immediate impact target sets such as energy and economic centers.
- **High value:** Contribute to degradation of societies ability to respond militarily or sustain itself economically. Targets include military, law enforcement and emergency response centers, Federal Government centers and critical commerce personalities.
- **Low value:** Contribute to localized fear and harassment of society and target sets include local transportation and non-critical infrastructure.

[7] Nance, Malcolm (2003). "The Terrorist Recognition Handbook," The Lyons Press, Gilford, Connecticut.

- **Tactical value:** Degrade local law enforcement capabilities to respond and includes target sets such as individual or small numbers of military or police, low level civil, military and law enforcement leadership personnel and centers, and military bases and equipment.
- **Symbolic value:** Heighten public fear and targets include innocent people, national treasures and landmarks, prominent public structures, and national representatives or diplomats.
- **Ecological value:** Damage natural resources of a society such as large bodies of natural resources and wide areas of agricultural resources and industry.

The terrorist target selection will likely be driven by the ultimate goal of its leadership, the feasibility of achieving success based on reports from the intelligence cells, and the ability to covertly deploy necessary cells to carry out the act.

Terrorist attack profiles are driven by the time to develop and execute a plan and they can use hard entry where they go in loud immediately with assaults using a range of weapons or use a soft (stealth) entry where penetration is not known until the attack occurs. They employ strategies that include misdirection (feints), deception (mask who or intent), and large numbers of identical incidents over a period of time. Planning and execution times can range from a few hours (hasty) to weeks (normal) to months and even years (deliberate). The terrorist methods and tactics vary. They have already demonstrated the use of hijackings, kidnappings, bombings, surface-to-air missiles, man portable air defense systems, arson, assassinations, armed assaults, and barricade-hostage incidents to attack critical infrastructure or capabilities, popular or high profile individuals, or important facilities or symbols. Weapons of mass effects (e.g., human suicide/martyr bombers, truck/car bombs, aviation attacks, maritime attacks, psychological, agriculture, ecological, economic, cyber) have been used as well and there is concern that they may in the future use weapons of mass destruction (e.g., chemical, biological, nuclear).

Initiation, escalation, de-escalation and termination of terrorist actions are determined by the leadership intent, the group capabilities (resources and expertise) and opportunities presented for attack. Terrorists have attacked both strategic and tactical targets worldwide — the intent is to make their presence felt. Western governments' security services have been reticent about sharing intelligence and judicial authorities rarely entertain request for extradition that adds to the difficulties of fighting the war on terrorism. Another important factor is the global media who are largely unaccountable to society and provide an unsophisticated form of terrorist Intelligence Surveillance and Reconnaissance, e.g., through transmission of live images of terrorism related events and by talking head analysis and special coverage assessments. Terrorists use symbolic acts to attract media and reach a large audience. They exploit the media to gain public attention, publicize their cause, and influence and spread fear. The media often make the mistake of seeking deeper goals in a terrorist operation than the terrorist set for them. This makes the terrorist appear powerful and untouchable. Media actions can also contribute to amplifying fear—a terrorist objective.

2.5 U.S. National Strategies

Following the terrorist attacks of September 11, 2001, the Bush administration developed and published seven national strategies that relate, in part or in whole, to combating terrorism and homeland security. These were:

- The *National Security Strategy of the United States of America*, September 2002.
- The *National Strategy for Homeland Security*, July 2002.
- The *National Strategy for Combating Terrorism*, February 2003.
- The *National Strategy to Combat Weapons of Mass Destruction*, December 2002.
- The *National Strategy for the Physical Protection of Critical Infrastructure and Key Assets*, February 2003.
- The *National Strategy to Secure Cyberspace*, February 2003.
- The *2002 National Money Laundering Strategy*, July 2002.

The *U.S. National Strategy for Combating Terrorism* is mainly offensive oriented but does include defensive homeland security objectives as well as objectives for protecting U.S. citizens abroad.[8] The principles of this strategy were used as a guide in the development of the case study counter terrorism influence net, scenarios and courses of action assessments discussed herein. The intent of the national strategy is to prevent, spoil actions, deter, and respond; neutralize or destroy terrorist groups; prevent attacks and minimize effects should one occur; weaken terrorist organizations and their political power; and make potential targets more difficult to attack. The goals and objectives of the 4D strategy (Defeat, Deny, Diminish and Defend) include:

Defeat terrorists and their organizations
- Attack sanctuaries, leadership, C3, logistics, and finances
- Disrupt ability to plan and operate
- Disperse and isolate terrorist
- Coordinate and use regional partners to neutralize terrorists

Deny further sponsorship, support and sanctuaries to terrorists
- End state sponsorship of terrorism
- Ensure regional states accept responsibilities to take action
- Interdict and disrupt material support for terrorist

Diminish the underlying conditions that terrorist seek to exploit
- Enlist international community to focus on areas most at risk
- Work with partners to keep combating terrorism
- Win the war of ideas

Defend U.S. citizens and interest at home and abroad
- Attain domain awareness
- Protect the homeland and extend our defenses to insure we identify and

[8] GAO-04-40ST. "Combating Terrorism: Evaluation of Selected Characteristics in National Strategies Related to Terrorism," February 2004.

neutralize the threat as early as possible

Success is dependent upon sustained, steadfast, and systematic application of all the elements of national power—diplomatic, economic, information, financial, law enforcement, intelligence, and military—simultaneously across all fronts.[9]

2.6 Terrorist Threat Considerations and Trends

The number of international terrorist attacks has declined but the level of violence and lethality has increased.[10] Primary sources of terrorist organizations are organized groups that have political, ethnic, and religious agendas; state sponsored organizations; transnational groups with broader goals; and Islamic terrorist groups that have become a growing threat. Al-Qaeda is gaining in global presence. These groups are loosely organized; recruit membership from many different countries; and obtain support from informal international networks. Terrorists have employed a wide variety of tactics to attack American targets worldwide that range from violent demonstrations to kidnapping to hostage taking to murder to armed attacks to bombings. Bombings are the most common type of attack (67% of all attacks against Americans).[11] Terrorist attack American businesses most frequently (more than 89% of the attacks) since businesses tend to be less protected and soft targets. U.S. government, diplomatic and military facilities tend to be protected and harder targets and less likely to be attacked. Terrorism varies by region of the world but most attacks occur in Latin America (87%).[12]

The reduced international barriers of the post-cold war landscape provide opportunities to exploit reduce political and economic barriers and facilitate movement of people, money, information and material across international borders. The global business networks facilitate international terrorism by providing safe havens for planning operations and allowing the terrorists to take advantage of global banking, communications, and transportation to carry out operations. Trafficking in narcotics, persons and weapons and organized crime are key sources of finance for operations.[13]

Other aggravating factors included technology advances and weak international law enforcement institutions. Information technology and communications facilitates global reach and terrorists are becoming more sophisticated in use of computer and telecommunications technology. Cell phones and Internet are used for planning, coordination, and execution. There are serious vulnerabilities in our critical infrastructure due to the reliance in information technology. The terrorists are adept at using technology for counterintelligence. Weak law enforcement institutions due to ineffective police and judicial systems in many foreign countries are a problem. Many of these institutions lack resources. There are outdated laws in many countries and some foreign governments are plagued by corruption. Law enforcement is constrained by national boundaries. Terrorists

[9] "National Strategy for Combating Terrorism," February 2003.
[10] GAO-03-165, "Combating Terrorism: Interagency Framework and Agency Programs to Address the Overseas Threat," May 2003.
[11] Ibid
[12] Ibid
[13] Ibid

take advantage of institutional limitations and weaknesses to find and establish sanctuaries.

Recent U.S. actions seem to have resulted in a decline in state-sponsorship of terrorism. Threats of sanctions and retaliation have reduced willingness of nations to support terrorist organizations. Terrorists have become less dependent on sponsorship by sovereign states and a new phenomenon is emerging—terrorist sponsoring a state (e.g., Taliban in Afghanistan). Terrorist groups operating on their own in loosely affiliated groups is on the increase as dependency on state sponsorship decreases. The terrorist organizations recruit membership from many different countries and obtain support form an informal network of like-minded extremists. There is a shift from aircraft hijacking and hostage taking to indiscriminate terrorist attacks that yield maximum destruction, casualties, and impact. This has generated a concern that there may be a shift to unconventional weapons of mass effects or even mass destruction. Alliances with transnational crime are providing the terrorist with access to various international crime organizations to help finance their operations.

2.7 Counter Terrorism Actions

The key to defeating terrorists lies in the realms of intelligence and police work, with military forces playing an important but nonetheless supporting role. Military destruction of al-Qaeda training and planning bases in Afghanistan have been successes in the war on terrorism but good intelligence—and luck—has formed the basis of virtually every other U.S. success against al-Qaeda.[14] Intelligence-based arrests and assassinations, not military divisions destroyed or ships sunk, are the cutting edge of successful counter terrorism actions. The war on terrorism is analogous to the international war on drugs. An effective strategy for counter terrorism needs to mobilize all elements of national power as well as the services of many other countries. Hence, to suppress, if not destroy, transnational terrorism it will be necessary to attack and destroy not their system architecture but their operational architecture–their ability to conduct operational activities in support of their goals.

There are numerous factors to consider as one builds a strategy for attacking the terrorist operational architecture. It is of utmost important to know your enemy in terms of motivation, his strengths and weaknesses, social networks of influence, sources of financing, logistics and other support, recruiting process, means of communicating, and organization structure and behavior. It is important to identify and locate terrorists and terrorist organizations then destroy them and their organizations. This requires an aggressive offensive strategy that aims to disrupt, dismantle, and destroy terrorist capabilities to carry out their operational activities by attacking their sanctuaries, leadership, C3I, material support, and finances.

The strategy needs to employ diplomatic, military and law enforcement means to eliminate sources of financing. As noted earlier, actions need to be taken to choke off the

[14] Record, Jeffrey (2003). "Bounding the Global War on Terrorism," Army War College, Strategic Studies Institute.

lifeblood of terrorist groups by employing the full range national power to end the state sponsorship of terrorism, to establish and maintain international accountability, to strengthen and sustain international effort to fight terrorism, to interdict and disrupt material support for terrorists, to eliminate terrorist sanctuaries, and eliminate conditions that terrorist can exploit.

Major threats to U.S. and world order today come from weak, collapsed, or failed states. Of concern is the fact that weak or absent government institutions in developing countries form the thread that links terrorism and weapons of mass destruction. Before 9/11, the U.S. viewed with less concern the chaos in far away places such as Afghanistan, but with the intersection of terrorism and weapons of mass destruction, these areas have become of major concern to the U.S. national security interests. Our tolerance for failed states has been reduced by the global war on terrorism and necessitates that we not leave weak and failed nations crumbling and ungoverned. Terrorists seek out such places to establish training camps, recruit new members, and tap into a black market where all kinds of weapons can be found for sale.[15] Courses of action to counter terrorism need strong consideration of ways to help rebuild and strengthen weak states and to identify and diminish conditions contributing to weak states by helping resolve poverty, deprivation, social disenfranchisement, and unresolved political and regional disputes. Partnering with the international community will be key. The strategy needs to win the war of ideas by employing actions that de-legitimize terrorism, kindle the hopes and aspirations of freedom, and support moderate and modern governments, especially in the Muslim world and in this regard assure Muslims that American values are not at odds with Islam. It will be necessary to reverse the spread of extremist ideology and to seek non-support, non-tolerance, and active opposition to terrorism from the international community. Use of effective, timely public diplomacy and government-supported media to promote the free flow of information and ideas will be needed as well.

The best defense is a good offense. This means investment of political will and resources to improve intelligence and warning and intelligence sharing among the military, law enforcement and our international partners. It will be necessary to integrate information sharing across the federal government and to effectively use intelligence, information and data across all agencies. Continuous law enforcement, intelligence and military pursuit of terrorists and their supporters will be necessary and needs to include a coordinated and focused effort of federal, state and local government, the private sector, and the American people. We will need to mobilize and organize to secure the homeland. In this regard, protection of vital systems and infrastructure is a shared responsibility of the public and private sectors. Plans need to be developed for alerting, containing and if necessary, repelling attacks. Measures to ensure the integrity, reliability, and availability of critical physical and information-based infrastructure at home and abroad need to be enhanced.

As noted earlier, intelligence is a key element of success in counter terrorism actions. The safe house is one of the key nodes of a terrorist operation and if seized may compromise cells, plans and materials. A safe house may be detected by informants, suspicious neighbors or through surveillance. Logistic cells have a higher probably of detection

[15] The Atlantic Monthly article "Nation Building 101" by Francis Fukuyama

because they often deal with low-level criminals and open market purchasing. Modern terrorist have become creative in the use of advanced information technology to conduct command and control of their operations making it difficult to detect activities. Terrorist can use diverse methods to finance their operations that include sources such as charitable organizations, organized crime, state sponsors, and legitimate business investments. Terrorist activity detection opportunities include:

- Leadership behavior
- State sponsors and other supporters
- Political and religious influence networks
- Safe houses
- Supply chains
- Logistics cells
- Storage of supplies
- Transportation and mobility
- Command, control, communications and intelligence
- Media relations and uses
- Financing
- Recruiting
- Training camps

The challenge to the intelligence and law enforcement community becomes one of asset management and focus and the ability to effectively share information and leverage the resources of the military, law enforcement and international community. A measure of success for a counter terrorism strategy will be diminished incidence and scope of terrorist attacks. However, analytically, this is an unsatisfactory measure of success since there is no way to prove a cause effect relationship. Additionally, a successful counter terrorism strategy can have self-defeating unintended consequences such as the terrorists changing their behavior and strategies that make them even harder to identify and neutralize. The GMU tool, CAESAR II/EB, may be of help to understand possible cause effect relationships of proposed courses of action and to identify potential unintended and undesired consequences. Successful results in this regard are highly depended upon the subject matter expert contributions and the creativity of the analyst constructing the influence net and the assessing the courses of action—it's an art not a science.

2.8 Counter Terrorism Case Study

The purpose of the case study was to demonstrate the utility and examine the challenges of using CAESAR II/EB to develop and assess EBO-based Courses of Action (COA) to mitigate an attack by a terrorist field cell by employing a broad-based strategic level attack profile that used both lethal and non-lethal means to disrupt and destroy the operational and systems architectures of the terrorist organization. The strategies tested employed the elements of National Power (Diplomatic, Information, Military, and Economic) to attack the terrorist organizations centers of gravity (Political, Religious, Military, Economic, Social, Infrastructure and Information). The study examined reactive, proactive, preemptive, and preventative tactics and examined the role of intelligence, the media, and the use of non-lethal means, such as, IO, Political, Legal, and International Collaboration. Homeland Security preparedness measures to defend high value targets was addressed as well.

Building the Model

Extensive research of the literature on historical experience with terrorism and strategies and frameworks for modeling counter terrorism actions was necessary in order to develop the understanding needed to create influence nets that could be used to assess counter terrorism courses of action and to examine the assessments for possible unintended consequences of actions taken against the terrorists and their organizations. Two RAND publications were of extreme value in the development of the case study influence net: the Paul Davis book titled "*Deterrence and Influence in Counter Terrorism*" and the Brian Jenkins book titled "*Countering al-Qaeda.*" A *Signal Magazine* article from the December 2001 issue by Dr Roger Smith, Titan Systems Corp., titled "*Counter Terrorism Modeling and Simulation: A New Type of Decision Support Tool*" was useful as well.

There are a number of interrelated challenges in constructing a counter terrorism influence net. First, is being able to think in terms of how the individuals and organizations to be modeled and attacked perceive they can be influenced and attacked—view the situation from the terrorist perspective. Second, is identifying the actors and the types and sequence of actions that can be taken to create the desired influence and behavior change. Additionally, thinking about whether terrorists and their organizations can be deterred, destroyed, or otherwise influenced requires a decomposition of the terrorist operations and supporting systems into classes of influence to be attacked.[16] Estimating the relative degree of impact of actions and events to influence outcomes needed to be developed and this proved to be a challenge as well—open literature documentation discusses the subject in qualitative terms.

The model for the case study was done at the strategic level and addressed broad-front national level actions needed to achieve an outcome that deterred a terrorist field cell from attacking. Past experiences using CAESAR II/EB to develop models in support of Naval War College Global war games[17] and Joint Forces Command experiment MC02[18] demonstrated that it was difficult to model at the operational level and much more difficult at the tactical level, and therefore, this effort focused on the strategic level.

The types of influence that needs to be considered can have both a positive and negative impact on the desired effect or event and determining the appropriate balance of these influences to achieve the desired effect is a challenge. It's largely a trial and error

[16] Davis, Paul and Jenkins, Brian (2002). "Deterrence and Influence in Counter Terrorism," RAND.

[17] Wagenhals, L. W. and Levis, A. H. (2002). "Modeling support of Effects Based Operations in War Games," 7th Command and Control Research and Development Symposium, Naval Post Graduate School, Monterey, CA, June 2002.

[18] Wentz, L. K. and Wagenhals, L.W. (2003). " Effects Based Information Operations," 8th International Command and Control Research and Technology symposium, National Defense University, Washington, D.C., June 2003.

experimentation process. For example, the higher the terrorist motivation and ability to attack, the less effective deterrence is likely to be. On the other hand, if the terrorist target of interest is well protected, the greater the deterrence. The influence net created for the case study is depicted in Figure 3 and was used to assess courses of action that reduced the probability that a terrorist field cell would attack.

The terrorist centers of gravity to be influenced and attack strategies ranged from using soft means to attack the political, social, belief, and financial structures to hard kill military means that disrupted or destroyed training facilities, logistics operations, weapons caches, and C3I capabilities needed to conduct operations. Threats to things terrorist care about, such as, loved ones, the terrorist cause itself, and the terrorist personal power and possessions are important deterrence factors and were the target of the IO campaign to influence perceptions, legal actions to seize possessions, and military and law enforcement actions to enforce messages in the IO campaign—actions need to support words. Other factors such as senior terrorist leadership support of terrorist cells and cause, continuation of state sponsorship of terrorists, continued approval by supporters of the terrorist and their cause, terrorist ability to conduct C3I of their operation, and the ability of the terrorist to finance operations are enablers and as such need to be attacked by an appropriate combination of all means available, especially the non-lethal means where and when possible. Public fear and anxiety are terrorist enablers that require careful attention and actions to keep the public informed and in this regard, both the government actions and the media messages play an important role in informing and influencing public understanding. Protection of high value targets is deterrence and this requires proactive government (federal, state and local) attention to protection policies, response plans and capabilities, and strategies and investments to protect critical infrastructure and key leadership personnel. Industry also has a role to play in investing in protection of facilities, capabilities, and key personnel. Awareness campaigns to educate and inform the public and make the terrorist aware that antiterrorism investments are being or have been made is important as well.

These considerations were built into the influence relationships and actions illustrated in the influence net shown in Figure 3. The desired outcome of the courses of action implemented is to drive the probability that a "terrorist field cell" will attack as low and as quickly as possible without creating unintended consequences such as windows of opportunity and vulnerabilities for the terrorist to attack. Key high level influence elements in the upper right hand quadrant of the influence net shown in Figure 3 include terrorist motivated to attack, finances available to conduct operations, recruiting and training capability providing new terrorist, the terrorist C3I capabilities able to support command and control of operations, logistics functioning and weapons available to support an attack, sanctuaries available to attack from, and continued approval of supporters such as political and religious leaders and other supporters of their cause exists. The lower right hand quadrant addresses perceptions of uncertainty and risk in terms of public fear and anxiety in response to terrorist threat warnings and terrorist belief that government and industry made the antiterrorism investments needed to protect high value targets (infrastructure—power, water, transportation—and leadership). The upper left hand quadrant includes influence elements such as state sponsorship, terrorist leadership support and media reporting of terrorist threats and terrorist perception of threats to things they care about.

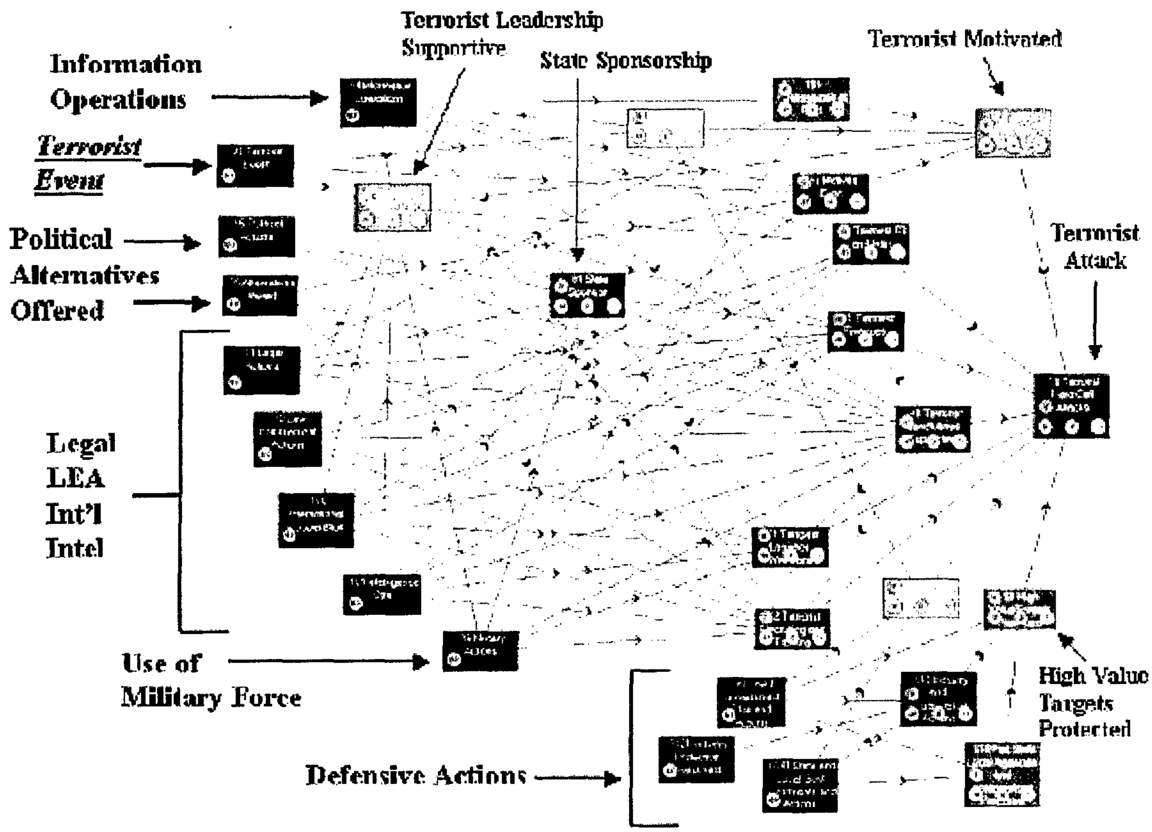

Figure 3. Counter Terrorism Influence Net

The left side of the influence net and lower right quadrant of the net show the different domains of actionable events. There are hard kill actions aimed at destroying terrorist targets that are largely military actions but law enforcement plays a role as well. The Intelligence action is the means to identify and monitor targets of opportunity, to develop social network understanding, to assess terrorist C2 tactics, procedures, and capabilities, and to develop situation awareness and actionable intelligence and warning. International cooperation is an enforcement enabler to provide an integrated global reach to leverage the use of other nations to help attack terrorist elements in their geographic area, to collect and share intelligence on terrorists, and to influence state sponsors and other supporters of terrorism to stop. Legal and law enforcement actions use international and national laws, law enforcement and judicial systems to disrupt terrorist organizations by arresting leadership and other members, disrupting terrorist recruiting activities, dismantling training camps, preventing cross border operations such as weapons trafficking and movement of terrorists, and dismantling of the terrorist financial networks. The political actions aim to gain international support to impose sanctions and to influence state sponsors and nations providing sanctuaries and other support to terrorist organizations and operations. The Information Operations actions focus on perception management of regional and local political and religious leaders, influencing the beliefs of the terrorist leaders, state sponsors, and members of the terrorist organizations and their supporters, and disruption of the recruiting of terrorists. An action referred to as "Alternatives Offered" aims to provide hope and improvements in quality of life of those suffering from poverty, deprivation and suppression of human rights who in turn support the terrorist cause and are a source of terrorist recruits. The provision of hope and improved the quality of life could serve to influence a large number of these people to quit supporting the terrorists and their cause. The lower right quadrant addresses federal, state and local government and industry actions (policies, contingency response plans, command, control and intelligence capabilities, and investments in infrastructure and key personnel protection) needed to implement antiterrorism measures to secure and protect high value targets and to be able to more effectively respond to indications of possible terrorist attacks.

The upper left hand quadrant has an action titled "terrorist event" and this was used as an intelligence and warning (I&W) indicator that a major terrorist attack was about to happen. Activation of this action served two purposes. First, its activation was used to positively influence the terrorist leadership support and motivation of the members of terrorist organizations and to influence the media response to generate radio and television public awareness messages and "talking head" discussions of the possibility and implications of an attack. The media response in turn had an additional positive influence on the motivation of the terrorists and in publicizing their cause. It also had a negative influence that contributed to the generation of public fear and anxiety.

A scenario-based approach was used to assess various courses of action so the second use of I&W actions was a trigger to initiate various courses of action strategies to be tested—reactive, proactive, preemptive, and preventative. In this role, the I&W action was used in two modes, the action could be turned on for the entire assessment timeframe or it could be turned on and off several times over the assessment timeframe to simulate multiple occurrences of threat warnings coming and going. The former mode was used to assess the impact of individual and various combinations of actions in response to the threat of a terrorist attack. The latter mode was used to assess the relative effectiveness of implementing course of action strategies that reacted to multiple warnings of terrorist attacks.

Sample COA Assessments

A number of assessments of the relative impact of individual and multiple actionable events on reducing the probability of attack and the sequencing and timing of these events were conducted as part of the research. Several different scenarios were also postulated based on the *U.S. National Strategy for Combating Terrorism* and used to formulate courses of action tested and assessed. Two examples of scenario-based courses of action assessments follow to illustrate the use of the tool and types of analysis conducted. The first example examines a strategy that reacts to multiple terrorist threat warnings and the second is a preemptive strategy in response to an initial threat warning and aims to minimize the probability of an attack as quick as possible given there will be a subsequent indication that an attack might occur. The two examples used different scenarios and sequencing and timing of the actionable events. The objective was not to select the optimum strategy and course of action or to imply one strategy was better than the other but to simply illustrate the use of the tool to conduct a comparison of these two strategies based on the probability of a terrorist attack over time and to provide some analysis of the relative effects of various courses of action.

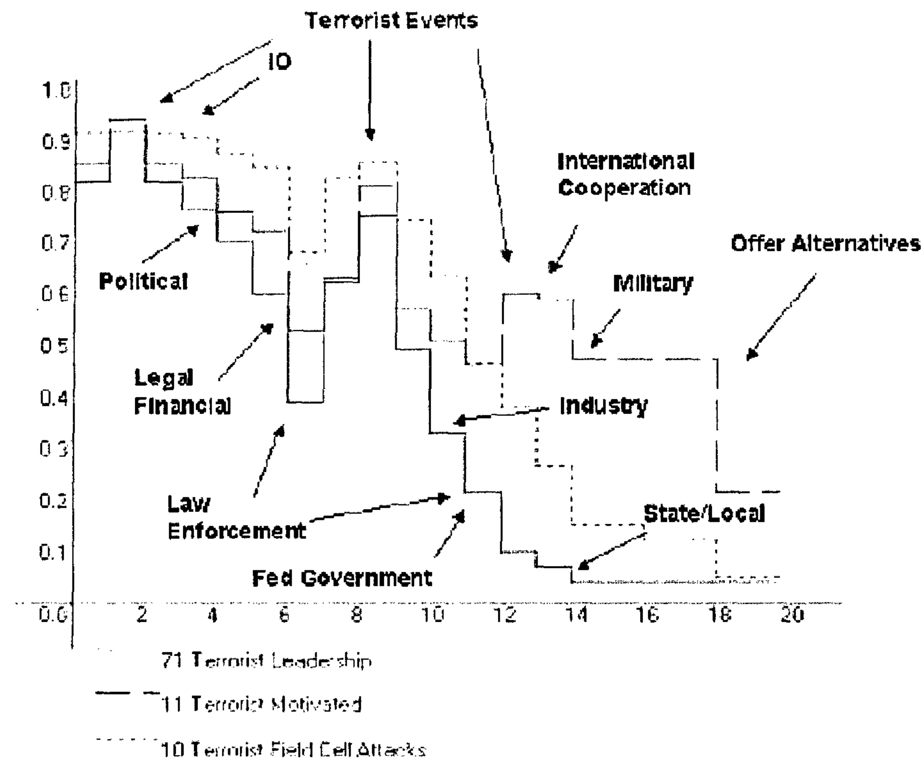

Figure 4. Reactive Strategy

The probability profiles in Figure 4 show the temporal analysis of terrorist leadership support, terrorist motivation and terrorist field cell likelihood to attack. Three single timeslot terrorist warning events occurred at times 1, 8 and 12 and these terrorist warning events were used to trigger a scenario-driven predetermined reactive course of action. The reaction strategy tested chose to use soft means first and then hard kill. IO followed by Political actions were initiated in reaction to the first terrorist threat warning event but these actions alone were not significant enough to cause a major reduction in the likelihood of a terrorist attack. The actionable events did serve to set some initial conditions for deterring an attack by reducing terrorist leadership willingness to support terrorist activities and there was some negative impact to terrorist motivation—largely driven by the IO campaign.

Legal and financial actions against the state sponsors and supporters and terrorist support elements such as sanctuaries and the financial networks were initiated at time 6. These actions combined with a short duration law enforcement action at time 6-7 against terrorist leadership and support elements appeared to have an important temporary impact on terrorist leadership, motivation and likelihood of attack. One might conclude that if the law enforcement action had continued (or its initial effects persisted) it would have helped reduce the relative influence of the second terrorist event that occurred at time 8. With the law enforcement action ending at time 7, it is suggested that a window of opportunity (or vulnerability) opened between time 7 and the next terrorist threat warning at time 8. As a result, the relative impact of the second threat warning was a more significant influence in raising the probability of attack.

Following the second terrorist event, the scenario proposed actions by government and industry to protect high value targets and this had a high payoff in reducing the probably of a terrorist attack. These actions increased the risk to the terrorists if they attacked. In this case, the scenario suggested that industry would respond quicker (loss of revenue driven) to the threats than government bureaucracies and that the federal government would be able to respond quicker than state and local governments and this drove the sequencing of antiterrorism protection actions. Law enforcement actions were also reactivated at time 11 to aggressively pursue terrorist leadership and support elements.

Although the third terrorist event increased terrorist motivation, the actions in place kept the probably that the terrorist would attack low—leveraged terrorist belief that attacking protected targets would be a high risk. Military action was initiated at time 14 to attack terrorist leadership and to reduce the ability of terrorists to conduct operations. The likelihood of a terrorist attack was further reduced when international cooperation and the offering of alternatives to improve the quality of life of terrorist supporters took place. These actions served to erode support for the terrorist cause and significantly reduced terrorist motivation.

Embedded within the temporal analysis shown in Figure 4 are multiple actions related to use of intelligence. The scenario assumed that there were limited intelligence assets available to support the counter terrorism and antiterrorism actions and that the use of these assets would therefore be driven by increased awareness that there was a need to focus on terrorism related targets. It was assumed that at time 0 that a minimum level of intelligence was being used (25%). Following the first terrorist event the use increased (50%) at time 5 but then went back down (25%) at time 7 when no attack occurred. Following the second terrorist warning event, the use was escalated (75%) and after the third warning its usage went to the max (100%).

The analysis suggests that an effective antiterrorism protection campaign can have a significant impact in reducing the likelihood of a terrorist attack but this alone is not sufficient. Other means need to be employed to dismantle the terrorist operational architecture—their ability to conduct operational activities in support of their goals. The scenario for the second example employed a preemptive strategy in response to a terrorist threat warning. In this case, proactive use of the elements of national power were brought to bear early with an aggressive combined use IO, intelligence, political, military, legal, financial and law enforcement actions to achieve an early deterrence in the probability of attack by going after the leadership, state sponsors, reducing terrorist motivation and disrupting their ability to conduct operations. The aggressive strategy was intended to buy time to allow the bureaucratic process to take the actions necessary to initiate protection of high value targets and to engage the cooperation of the international community that would in turn serve to reduce the likelihood of an attack by further reductions in state sponsorship, terrorist supporters and support activities and the elimination of sanctuaries.

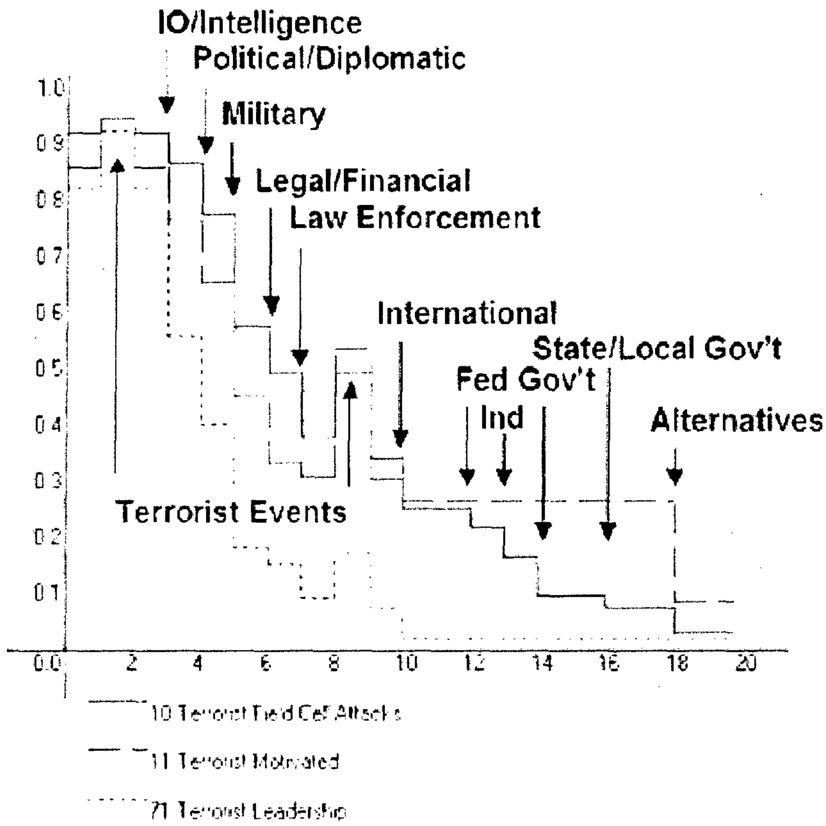

Figure 5. Preemptive Strategy

The probability profiles in Figure 4 show the temporal analysis of terrorist leadership support, terrorist motivation and terrorist field cell likelihood to attack. There are two terrorist warning events, one at time 1 and a second at time 8. The first terrorist warning event triggered the response to aggressively attack. The resulting effect was to drive the probability that the terrorist would attack below 50% and even with the second attack warning the probably of attack did not rise above 50%. Initiation of international cooperation at time 10 served to further reduce terrorist leadership willingness to support terrorist attack actions and this influenced a reduction in terrorist motivation and willingness to attack. As was the case in the first example, initiation of antiterrorism protection actions caused a reduction in the likelihood of a terrorist attack and the offering of alternatives to improve the quality of life of terrorist supporters served to further reduce terrorist motivation. The results suggest that the aggressive attack strategy was successful in achieving an early dismantling of the terrorist ability to conduct operations and significantly reduced the leadership support and other support of terrorist actions. Although the probability that the terrorist would attack was driven below 50% before the second terrorist warning event, the results also suggest that an aggressive antiterrorism program is needed to compliment the aggressive counter terrorism program. Both examples suggest that neither alone is sufficient.

The CAESAR II/EB tool has an ability to do a sensitivity analysis of the relative impacts of individual and combinations of actions. A sensitivity analysis of the case study model suggested that international cooperation and IO were key actions that if used in combination with other lethal and non-lethal actions could be a force multiplier and important contributor to reducing the probably the terrorist cell would attack.

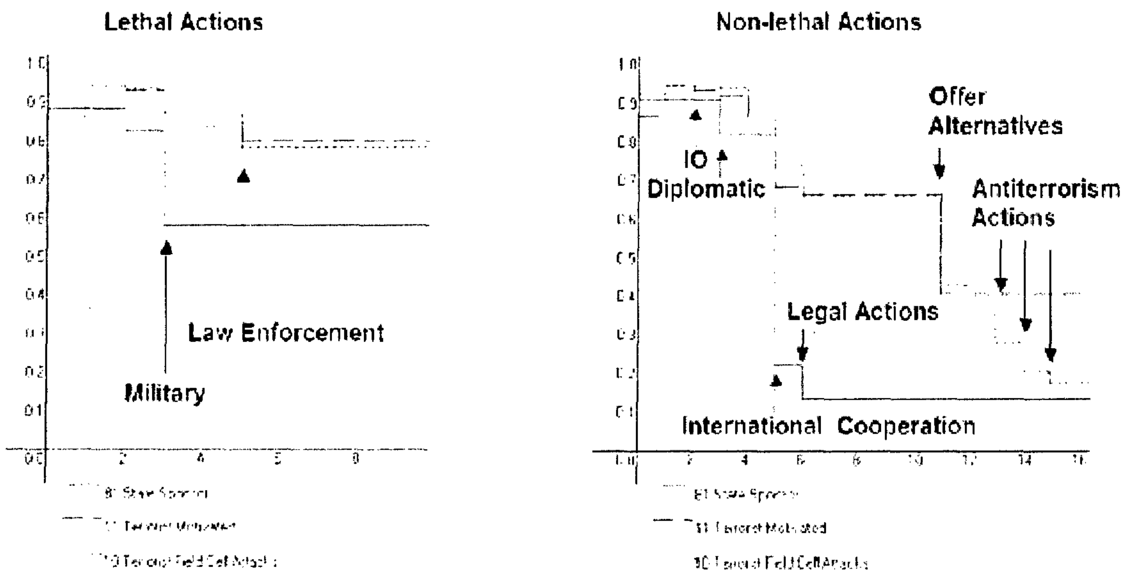

Figure 6. Comparison of use of Lethal and Non-lethal Means

Figure 6 compares the use of lethal and non-lethal means in response to the belief a terrorist event might occur. The probability profiles show the temporal analysis of state sponsorship, terrorist motivation and terrorist field cell likelihood of attack. In both cases, a terrorist warning event occurs at time 1 and at time 2 intelligence actions were initiated in response to this warning. The comparison suggests that although the follow on military and law enforcement actions reduced the willingness of state sponsors to support terrorist activities, these actions alone were not sufficient to significantly impact the terrorist willingness to attack. On the other hand, the use of non-lethal means such as IO, political/diplomatic, international cooperation and legal/financial actions appeared to be significantly more effective in terms of reducing state sponsorship willingness to support the terrorist activities but here too these alone were not sufficient to significantly reduce the probability the terrorist might attack. The follow on offering of alternatives to improve the quality of life of terrorist supporters drove the terrorist motivation down and the probability of attack below 50%. The antiterrorism protective actions served to further reduce the likelihood of attack—increased risk to terrorist but not a de-motivation of support of the cause. One might conclude from this assessment that non-lethal means can be a significant contributor to reducing the probability that the terrorist might attack. Comparing these results with the preemptive strategy illustrated in Figure 5 also suggests that combining early military and law enforcement actions with non-lethal means such as

IO, political, international cooperation and legal actions provided a synergistic effect (i.e., non-lethal means can be force multipliers) that achieved an early dismantling of the terrorist ability to conduct operations and reduced the willingness of the supporters to continue their support of terrorist actions and hence, satisfied the end objective to significantly reduce the probability that the terrorist would attack.

2.9 Observations

As noted earlier, creating influence nets and assessing courses of action is an art not a science. As such, the experience of the model builder is key as well as availability of subject matter experts to help guide the development of the models, the selection of courses of action and subsequent assessments. In many cases, the subject matter experts are not readily available and the modeler needs to do the research to prepare to develop the influence nets and conduct the course of action planning and assessments. This was the situation for the Counter Terrorism case study presented herein—a large part of the effort was researching the subject area. Model building is also a timely and complex task. In the authors' view, the current tools work best at the strategic level and to a limited extent at the operational level. The pace of tactical operations coupled with the author's experience using and observing the use of such tools in exercises and experiments suggests that these tools can be cumbersome to use operationally and hence, limit their value added in the high OPTEMPO environment of the tactical level of operation.[19]

The value added of CAESAR II/EB was successfully demonstrated at the strategic level when it was used to support the Naval War College Global Wargames and at the operational level when it was used to support the Joint Task Force Information Operations cell at the Millennium Challenge 2002 experiment at JFCOM. It must be remembered, however, that tools, such as CAESAR II/EB, are research tools and not ready for prime time operational use. Hence, the man-machine interfaces are not that user friendly and visualization of the results have limitations—CAESAR II/EB is cumbersome to use and generates probability profiles as its visualization output. Results must also be used carefully since this is just one means for trying to gain insights into effects actions might have on achieving a desired outcome. It's a prediction with varying degrees of uncertainty.

Challenges related to constructing influence nets are numerous. Understanding the situation is key to identifying the effects to be modeled and to develop the causal relationships and predict the truth or falsity of parent node effects on the child nodes. Selections of actions and the timing of the sequencing of these actions require some creativity on the part of the modeler as well. The process usually is to build a little and test a little with lots of trial and error experimentation to refine the model and to develop and select courses of action to be assessed. Models have limitations as well. For example, for CAESAR II/EB, persistence or the continuation of the effect after the action is removed is not modeled. Actions can be turned on and off several times over time but the

[19] Wentz, L. K. and Wagenhals, L.W. (2003). " Effects Based Information Operations," 8th International Command and Control Research and Technology symposium, National Defense University, Washington, D.C., June 2003.

persistence factor is not modeled. The model does not differentiate between the effects of the sequencing of two actions (e.g., action A before B versus B before A gives same final result although intermediate probabilities may be quite different) that in a real life situation may not be the case. On the other hand, the insights and interchanges among the decision makers, analysts and planners and synergy derived from the process of developing models and assessing the courses of action is probably one of the most important benefits to be realized from using a tool such as CAESAR II/EB.

The Counter Terrorism model developed using CAESAR II/EB and related courses of action planning and assessments appear to provide useful insights into the effects of lethal and non-lethal actions and their timing on desired deterrence outcomes as well as to help identify unintended and undesirable consequences of actions taken. The analysis presented herein suggests that counter terrorism and antiterrorism strategies need to address both the operational and technical architectures of the terrorist operations and organizations as well as one's own architectures. The experience has enabled the GMU researchers to expand their repertoire of modeling types and techniques to provide support to different classes of problems. CAESAR II/EB has limitations and work is in progress at GMU to explore enhancements to the utility of the tool including incorporation of modeling persistence and improving the user friendliness and visualization of results in support of effects based COA planning and assessments. Similar research and modeling efforts at the Air Force Rome Labs have already addressed some of these short falls. Their Causal Analysis Tool has incorporated modeling persistence and improved user interfaces and visualization and additional research is addressing improvements to the operational utility of CAT to support effects based air operations planning and assessments.

References

Anonymous (2003). "Terrorist Hunter," Harper Collins Publisher.

Cordesman, Anthony (2001). "A new US Strategy for Counter Terrorism and Asymetric Warfare," Center for Strategic and International Studies, Washington, D.C.

Davis, Paul and Jenkins, Brian (2002). "Deterrence and Influence in Counter Terrorism," RAND.

Ehrenfeld, Rachel (2003). "Funding Evil," Bonus Books.

Gunaratna, Rohan (2002). "Inside al Qaeda," Berkley Books.

Hess, Stephen and Kalb, Marvin (2003). "The Media and the War on Terrorism," Brookings Press.

Hoffman, Bruce (1998). "Inside Terrorism," Columbia University Press.

Hudson, Rex (1999). "Who Becomes a Terrorist and Why," The Lyons Press, Gilford, Connecticut.

Internet. "Terrorism Research center," www.terrorism.com

Jenkins, Brian (2002). "Countering al Qaeda," RAND.

Ledeen, Michael (2003). "The War Against the Terror Masters," Truman Talley Books.

Record, Jeffrey (2003). "Bounding the Global War on Terrorism," Army War College, Strategic Studies Institute.

Lesser, Ian and et al (1999). "Countering the New Terrorism," RAND.

Nance, Malcolm (2003). "The Terrorist Recognition Handbook," The Lyons Press, Gilford, Connecticut.

Norris, Pippa and et al (2003). "Framing Terrorism," Routledge Press

Palmer, Nancy (2003). "Terrorism, War, and the Press," Harvard University.

Wagenhals, L. W. and Levis, A. H. (2002). "Modeling support of Effects Based Operations in War Games," 7th Command and Control Research and Development Symposium, Naval Post Graduate School, Monterey, CA, June 2002.

Wentz, L. K. and Wagenhals, L.W. (2003). " Effects Based Information Operations," 8th International Command and Control Research and Technology symposium, National Defense University, Washington, D.C., June 2003.

SECTION 3

Transforming Timed Influence Nets into Time Sliced Bayesian Networks

Sajjad Haider and Abbas K. Zaidi

3.1 Introduction

The easy access to domain-specific information and cost-effective availability of high computational power have changed the way people think about complex decision problems in almost all areas of application, ranging from financial markets to regional and global politics. These decision problems often require modeling of informal, uncertain, and unstructured domains in order for a decision maker to evaluate alternates and available courses of actions. The past few years have witnessed an emergence of several modeling and analysis formalisms that try to address this need. The modeling of an uncertain domain using Probabilistic Belief Networks, or more commonly known as Bayesian Networks (BNs), is considered to be the most used and popular of all such formalisms. The BN approach requires a subject expert to model the parameters of the domain □ random variables □ as nodes in a network. The arcs (or directed edges) in the network represent the direct dependency relationships between the random variables. The arrows on the edges depict the *direction* of the dependencies. The strengths of these dependencies are captured as conditional probabilities associated with the connected nodes in a network. A complete BN model requires specification of all conditional probabilities prior to its use. The number of conditional probabilities on a node in a BN grows exponentially with the number of inputs to the node. The requirement of specifying an exponentially large number of conditional probabilities presents a, at times insurmountable, modeling challenge. Cheng et al. [1994] developed a formalism, at George Mason University, called CAusal STrength (CAST) logic, as an intuitive, and approximate language to elicit the large number of conditional probabilities from a small set of user-defined parameters. The logic requires only a pair of parameter values for each dependency relationship between any two random variables. The CAST logic is used as a knowledge elicitation interface to an underlying BN. The approach was subsequently named Influence Nets [Rosen and Smith, 1996]. The Influence Nets require a system modeler (or subject expert) to specify the CAST logic parameters instead of the probabilities. The required probabilities are internally generated by the CAST logic with the help of user-defined parameters. The Influence Nets are, therefore, appropriate for modeling situations in which it is difficult to fully specify all conditional probability values and/or the estimates of conditional probabilities are subjective and estimates for the conditional probabilities cannot be obtained from empirical data, e.g., when modeling potential human reactions and beliefs.

Both Bayesian Networks and Influence Nets are designed to capture *static* interdependencies among variables in a system. A situation where the impact of a variable takes some *time* to reach the affected variable(s) cannot be modeled by either of the two approaches. In the last several years, efforts have been made to integrate the notion of time and uncertainty. Wagenhals et al. [Wagenhals et al. 1998] have added a special set of temporal constructs to the basic formalism of Influence Nets. The Influence Nets with these additional temporal constructs are called Timed Influence Nets (TINs). TINs have been experimentally used in the area of Effects Based Operations (EBOs) for evaluating alternate courses of actions and their effectiveness to mission objectives. The provision of time allows for the construction of alternate courses of actions as timed sequences of actions or actionable events represented by nodes in a TIN [Wagenhals and Levis, 2000; Wagenhals and Levis, 2001; Wagenhals et al., 2003].

The TIN approach inherits both the advantages and disadvantages of the Influence Net formalism: it offers an intuitive, and compact, knowledge elicitation interface for modeling purposes, but lacks some important analysis techniques. Currently, the analysis suite of TINs lacks the ability to incorporate the real-time information/evidence coming from different sources during the execution of a previously selected course of action. In a military/political scenario, this new information might come from the surveillance system regarding an adversary's actions. In an economic domain, a new development in a stock market, e.g., bankruptcy filed by some corporation, might be taken into account before making a strategic decision. In any case, this new information results in the revision of a previously held belief about some variables in the system. Haider and Levis [Haider and Levis, 2004] have recently proposed an algorithm to overcome this limitation; however, the approach is applicable for a special class of evidences only.

On a parallel track, scholars in the BN community extended the BN formalism to incorporate a special notion of time in it. The extension, called Time Sliced Bayesian Networks (TSBN) or Dynamic Bayesian Networks [Murphy, 2002], has gained a privileged status among the Artificial Intelligence community as a tool for modeling time and uncertainty. The approach is based on *unrolling* a static BN on a discrete time line with each time *slice* having an instance of a node in the network. The temporal dependencies are modeled with the help of edges across these time *slices*. Several sophisticated techniques for enhancing the capabilities of this approach have been proposed [Hanks et al., 1995; Figueroa and Sucar, 1999; Santos and Young, 1999; Galan and Diez, 2002]. Furthermore, several algorithms have also been proposed to compute the marginal probabilities of the random variables in an efficient manner [Kjaerulff, 1992; Boyen and Koller, 1998; Doucet et al., 2000; Murphy and Weiss, 2001; and Takikawa et al., 2002].

The lack of a comprehensive suite of analysis techniques in the TIN formalism and the recent developments in the field of TSBN bring us to the topic of this paper: The paper is an attempt to combine the advantages of both paradigms, TIN and TSBN, into a single formalism by providing a mapping from a TIN to a TSBN. The proposed formalism uses the TIN approach for the model building and the TSBN for analysis and evaluation. The paper demonstrates that TINs provide a compact and an intuitive way of modeling dynamic domains. A system modeler, therefore, can specify the uncertainties and the temporal constraints, present in a problem domain, in the form of a TIN. Once a TIN is fully specified, it can be converted into a TSBN using the approach presented in this paper. On one hand, the conversion simplifies the intractable task of knowledge elicitation in TSBNs by suggesting the use of TINs as a front end tool; while on the other, the conversion makes it possible to use a variety of analysis algorithms that have been developed for TSBNs.

The rest of the paper is organized as follows: Section 2 provides a technical background of Timed Influence Nets and Time Sliced Bayesian Networks. The algorithm for transforming TIN into TSBN is described in Section 3 with the help of examples. Finally, Section 4 discusses the conclusions and proposes directions for future research.

3.2. Technical Background

3.2.1 *Bayesian Networks*

Over the last two decades, Bayesian Networks, (BNs) have become a popular way of modeling uncertainty in several fields of studies [Pearl, 1987; Charniak, 1991; Jensen, 2001; Neapolitan, 2003]. A BN is a Directed Acyclic Graph (DAG) **G** = (**V**, **E**). The nodes or vertices (**V**) in the graph represent random variables while edges (**E**) connecting pairs of variables represent probabilistic dependencies between them. Definitions 2.1-2.3 present a formal description of BNs and the related terminology.

Definition 2.1 [Neapolitan, 2003]

Given a DAG **G** = (**V**, **E**) and nodes X and Y in **V**, Y is called a *parent* of X if there is an edge from Y to X, Y is called a *descendent* of X and X is called *ancestor* of Y if there is a path from X to Y, and Y is called a *nondescendent* of X if Y is not a descendant of X.

Definition 2.2 [Neapolitan, 2003]

Suppose we have a joint probability distribution **P** of the random variables in some set **V** and a DAG **G** = (**V**, **E**). We say that (**G**, **P**) satisfies the *Markov Condition* if for each variable X ∈ **V**, {X} is conditionally independent of the set of all its nondescendent given the set of all its parents.

Definition 2.3

Let **G** = (**V**, **E**) be a DAG and **P** be the joint probability distribution of **V**. If (**G**, **P**) satisfies the *Markov Condition* then **B** = (**V**, **E**, **P**) is called a *Bayesian Network* and **P** can be written as

$$P(x_1 \ldots\ldots x_n) = \prod_{i=1}^{n} P(x_i \mid pa(x_i))$$

where *pa(x)* represents the set of all parents of *x* and *x* ∈ V.

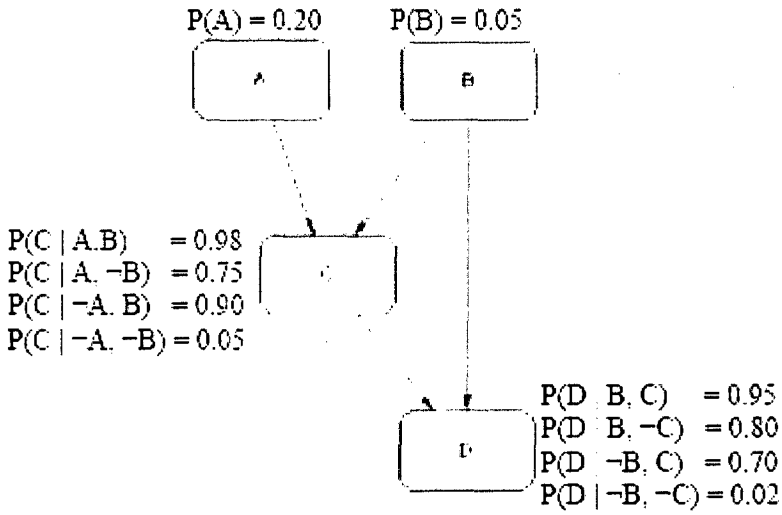

Figure 1: A Sample Bayesian Network

Figure 1 shows an example of a BN having four binary variables, namely, A, B, C, and D. The text in the figure shows the prior and conditional probabilities associated with the root and non-root nodes, respectively. The joint distribution of all the variables is computed as the product of these probabilities.

For example, P (A, ¬B, C, D) is computed as
P (A, ¬B, C, D) = P (D | ¬B, C) P (C | A, ¬B) P (A) P (¬B)

The other values in the joint distribution can be computed in a similar fashion. Once the computation of joint distribution is completed, it can be used to determine the marginal probabilities of the variables of interest. Several algorithms have been developed for various graphical structures of BN that compute the marginal probabilities in an efficient way by *propagating* probabilities without first calculating all the joint distribution values. The random variables in a BN could be either discrete or continuous. Binary random variables are considered a special case of discrete random variables. The approach in this sequel, assumes that all the variables in a BN have binary states. The presented approach can be extended to more general cases.

3.2.2 Time Sliced Bayesian Networks

A TSBN works by discretizing time and creating instances of variables in a BN for each point in the time interval under consideration. The process starts with the identification of static cause and effect relationships among the variables and then by repeating the same structure for multiple time slices. Links are drawn between variables having temporal dependencies. Suppose, in the model of Figure 1, the probabilities of node A and B at time t depend upon their probabilities at time t-1. Then, the probabilities of A and B at time 1 are influenced by their respective probabilities at time 0; the probabilities at time 2 are influenced by the probabilities at time 1, and so on. These temporal dependencies can be captured in a TSBN as shown in Figure 2.

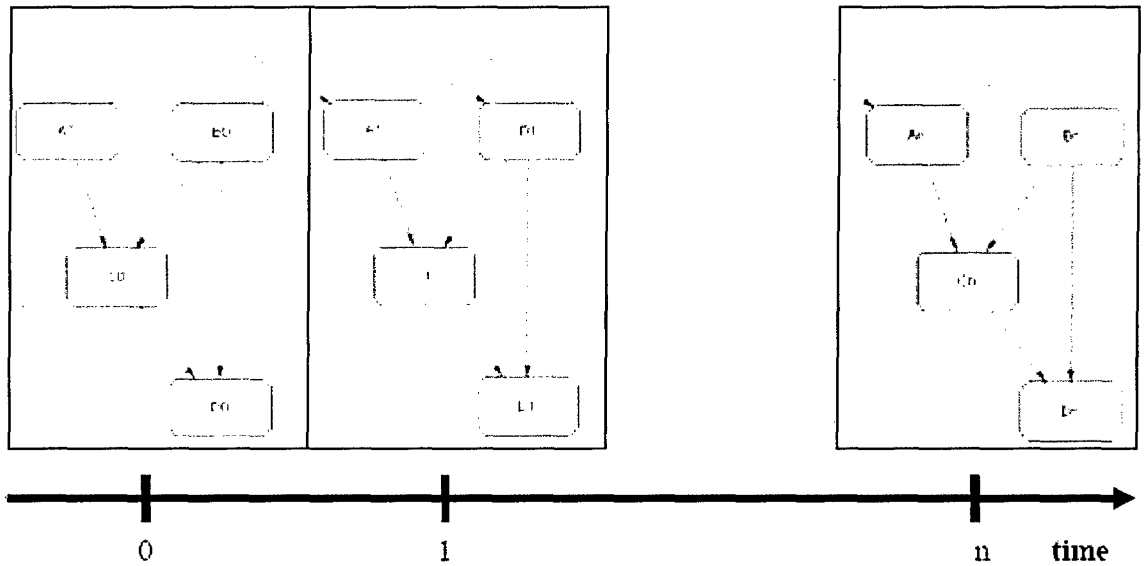

Figure 2: A Time Sliced Bayesian Network

Depending upon the situation, the state of a variable at time slice t may depend upon the states of influencing variables in preceding time slices ranging from t-1 to 0, i.e., the conditional probability of X at time t depends upon the set of influencing variable (the parents of X) at time slices ({t-1} or {t-1, t-2} or ....or {t-1,t-2,....,t-k}). In this 7 context, a TSBN can be seen as an order 'k' *Markov Chain.* Typically, the TSBNs are built as order one Markov Chain, i.e., the future is conditionally independent of the past given the present. One more assumption that simplifies the specification of a TSBN is that the changes in the state of variables are caused by *stationary processes.* In other words, it is assumed that the conditional probabilities do not change over time, i.e., $P(x_t \mid pa(x_t)) = P(x_{t-1} \mid pa\ (x_{t-1}))\ t = 1, \ldots, n$ In the sequel, a TSBN is assumed to have stationary conditional probabilities and of order one Markov process, unless stated otherwise.

3.2.3 Timed Influence Nets

As mentioned earlier, Influence Nets simplify the intractable task of eliciting Conditional Probability Tables (CPTs) from subject experts, especially when a node in the net has many parents. They use CAST Logic as an interface for eliciting CPTs. The logic has its origin in 'Noisy-OR' approach [Agosta, 1991; Drudzel and Henrion, 1993; Heckerman and Breese, 1996]. The CAST logic not only simplifies the elicitation of CPTs, but it also provides a mechanism to obtain information from various experts and then combine their individual assessments in a mathematical manner. The exact details of the CAST logic algorithm are beyond the scope of this paper. The interested reader should refer to Chang et al. [1994] and Rosen and Smith [1996].

Timed Influence Nets extend the capabilities of Influence Nets by allowing the provision of specifying several types of temporal information. These types can be broadly classified into two categories. One is related to the delays present in a problem domain while the other is related to the actionable events. The delays present in the domain represent the amount of time it takes for knowledge about a change, in the status of any variable, to be propagated to the node that is affected by that change. In TINs, this phenomenon is modeled by associating delays to arcs and nodes. The delay on an arc represents the communication delay, while a delay on a node represents the information processing delay. The second type of temporal information in TINs is associated with the actions taken in a course of action. Wagenhals et al. [Wagenhals et al., 2003] have called this type an *input scenario*. It describes the time at which the actions are taken and the intervals during which these actions are maintained. Actions in this context refer to the random variables that are modeled as root nodes in the corresponding TIN. In Bayesian literature, these actions could correspond to having the evidence on the root nodes. It is assumed that the actions occur instantaneously. Because of the dynamic nature of the problem, it is possible that the state of an action is changed during a course of actions. Thus, an action can be true during a particular time interval and false in another. Furthermore, these actions can be repeated an arbitrary number of times. It should be mentioned that causal strengths in TINs do not change over time. It is, therefore, assumed that like TSBNs, the changes in the state of variable in TIN are caused by stationary processes. The following items characterize a TIN:

1. A set of random variables that makes up the nodes of a TIN. All the variables in the TIN have binary states.
2. A set of directed links that connect pairs of nodes.
3. Each link has associated with it a pair of CAST Logic parameters that shows the causal strength of the link (usually denoted as g and h values).
4. Each non-root node has an associated CAST Logic parameter (denoted as baseline probability), while a prior probability is associated with each root node.
5. Each link has a corresponding delay d (where $d \geq 0$) that represents the communication delay.
6. Each node has a corresponding delay e (where $e \geq 0$) that represents the information processing delay.
7. A pair (p, t) for each root node, where p is a list of real numbers representing probability values. For each probability value, a corresponding time interval is defined in t. In general, (p, t) is defined as

$$([p_1, p_2, \ldots, p_n], [[t_{11}, t_{12}], [t_{21}, t_{22}], \ldots, [t_{n1}, t_{n2}]]$$
$$\text{where } t_{i1} < t_{i2} \text{ and } t_{ij} > 0 \quad \forall i = 1, 2, \ldots, n \text{ and } j = 1, 2$$

Formally, a TIN is described by either one of the following definitions (Defs. 2.4a, b, c).

Definition 3.2.4a

A *Timed Influence Net* is a tuple **(V, E, C, B, $\mathbf{D_E}$, $\mathbf{D_V}$, A)** where
V: set of Nodes,
E: set of Edges,
C represents causal strengths: $\mathbf{E} \rightarrow \{\, (\mathbf{h}, \mathbf{g}) \text{ such that } -1 \leq \mathbf{h}, \mathbf{g} \leq 1 \,\}$,
B represents Baseline / Prior probability: $\mathbf{V} \rightarrow [0,1]$,
$\mathbf{D_E}$ represents Delays on Edges: $\mathbf{E} \rightarrow N$,

$\mathbf{D_V}$ represents Delays on Nodes: $\mathbf{V} \to \mathrm{N}$, and
A (input scenario) represents the probabilities associated with the state of actions and the time associated with them.

$$\mathbf{A}: \mathbf{R} \to \{([p_1, p_2, \ldots, p_n], [[t_{11}, t_{12}], [t_{21}, t_{22}], \ldots., [t_{n1}, t_{n2}]]) \text{ such that } p_i = [0, 1],$$
$$t_{ij} \to Z \text{ and } t_{i1} \leq t_{i2}, \forall i = 1, 2, \ldots., n \text{ and } j = 1, 2 \text{ where } \mathbf{R} \subset \mathbf{V} \}$$

Definition 3.2.4a can be further simplified by reducing some of the elements in the tuple. The elements **C, B** in the tuple are used to approximate conditional probabilities, which in turn, are used to represent the joint distribution **P** of the random variables in **V**.

Definition 3.2.4b

Given a TIN ($\mathbf{V, E, C, B, D_E, D_V, A}$), the elements **C, B** can be replaced by **P** that represents the joint distribution of the variables in **V**. The transformation is done by CAST Logic.

$$\text{TIN} = (\mathbf{V, E, C, B, D_E, D_V, A}) \to \quad \text{TIN} = (\mathbf{V, E, P, D_E, D_V, A})$$

The elements $\mathbf{D_E}$ and $\mathbf{D_V}$ in the Definitions 3.2.4a and 3.2.4b represent the delays associated with the edges and nodes, respectively. The delay associated with a node can be remodeled by adding it to the delays on its incoming arcs and removing it from the corresponding node. For example, consider the TIN shown in Figure 3 (a).

(a) (b

Figure 3: Reassignment of Node Delays

The delays on the links between A-C and B-C are 2 and 1 time units, respectively. The delay associated with the node C is 2 time units. Figure 3(b) shows the equivalent net with the node delay transformed to the edge delays. This transformation yields the following definition for TINs.

Definition 3.2.4c

Given a TIN = ($\mathbf{V, E, P, D_E, D_V, A}$), the elements $\mathbf{D_E}$ and $\mathbf{D_V}$ can be mapped into an equivalent **D** that represents the transformed delays associated with the edges **E** in a TIN.

$$\text{TIN} = (\mathbf{V, E, P, D_E, D_V, A}) \to \text{TIN} = (\mathbf{V, E, P, D, A})$$

3.2 Transformation from TINs to TSBNs

The existing algorithms for TINs propagate the influence of actions in the forward direction only, i.e., the probabilities are propagated from source (input nodes) to sink (or target nodes) through intermediate nodes. This presents an analysis and computational limitation of TINs for situations where observations, regarding states of non-root nodes in a TIN, arrive during the execution of a selected course of action. An approximation algorithm [Haider and Levis, 2004] has been proposed for incorporating such observations. The algorithm, however, puts certain restrictions on the timing of input evidence, thus making it impractical for some cases. One way of overcoming this limitation is to transform a TIN into a TSBN. The exact details of the transformation algorithm are presented in this section.

3.3.1 The Algorithm

The transformation algorithm first determines the required number of time slices by taking the maximum length of the paths that exist between the root nodes and the target node. More slices are added later when a course of action is selected. The additional slices are determined by looking at the largest time stamp associated with the actions in the selected course of action. Later, the connections between the nodes are established based on the time delays associated with the arcs that connect two nodes in the TIN. The subsequent indices of a root node (representing actionable events) are also connected except for the time when an action is taken. The exact algorithm is presented in Table 1.

3.3.2 Application of the Transformation Algorithm

This section illustrates the transformation algorithm with the help of an example. The example model is shown in Figure 4. The TIN in the figure shows a hypothetical crisis that arose on a piece of land, belonging to a peaceful country, but annexed by a hostile country B. The objective of building this model is to explore the possibility of a peaceful solution of the crisis, or, in other words, the objective is to determine the probability that country B would agree to withdraw its forces based upon certain actions taken by the international community. There are four nodes in the TIN, namely, A, B, C, and D. The description of these nodes is shown in the figure. The text besides the links represents the delays associated with them. For instance, the link between B and C has a delay of 1 time unit while the link between C and D has a delay of 3 time units, and so on. The steps involved in the transformation algorithm are shown in Figure 5.

<u>Step 1:</u> In this step, the maximum path length between the root nodes (A and B) and the target node (D) is determined. The path A-C-D has the maximum length of 5. Thus, M is set to 5.

Table 1: The Transformation Algorithm

Given TIN = (**V**, **E**, **P**, **D**, **A**)

1. Find the maximum path length between the root nodes and target nodes, i.e.,

   $M = \max_{i,j}[P_{i,j}]$ where

   $P_{i,j}$: path between nodes i and j such that $i, k \in V$ and $\neg\exists\,(k, i) \in E$

2. Construct a TSBN (**V1**, **E1**, **P1**) where

   V1: $\forall\, v \in$ **V** add v_i to **V1** where $i = 0, 1, \ldots, M$

   $= \{v_i \mid v \in \mathbf{V}, i = 0, 1, \ldots, M\}$

   E1 $= \{(x_i, y_j) \mid i = \max(0, j - D(x, y));\ x, y \in \mathbf{V}$ and $i, j = 0, 1, \ldots, M\}$

   P1: **P** when indices are ignored

   For example, $P(y_i \mid x_i) = P(y \mid x)$ when $x, y \in \mathbf{V}$ and $x_i, y_j \in \mathbf{V1}$.

This step draws the nodes in the TSBN for M time slices. The connections are drawn between the non-root nodes and their parents. The following step is required once an input scenario is determined.

3. Let S = maximum time stamp associated with the root nodes as provided by the input scenario:

   (a) Add S additional time slices in the TSBN obtained in the previous step by following the procedure outlined in Step 2.

   (b) The resultant network is the modified TSBN (**V1**, **E1**, **P1**) where

   V1 $= \{v_i \mid v \in \mathbf{V}, i = 0, 1, \ldots, M+S\}$

   E1 $= \{(x_i, y_j) \mid i = \max(0, j - D(x, y));\ x, y \in \mathbf{V}$ and $i, j = 0, 1, \ldots, M+S\}$

   P1: **P** when indices are ignored

   (c) Let **R1** = Set of Root Nodes where **R1** $\subset$ **V1**. $\forall\, r \in$ **R1** connect r_{t-1} to r_t where $t = 1, 2, \ldots, M+S$, unless t is the time at which the variable is set to a state.

<u>Step 2:</u> This step draws the nodes in the TIN for M time slices in the corresponding TSBN. The step is shown in Figure 5(a). After drawing the nodes for 5 time slices, the connections between the nodes are drawn. The delays on the arcs in the TIN determine the indices of the connected nodes in the corresponding TSBN. For instance, the delay between B and D is 2. The connections between instances of B and instances of D are determined as shown below:

D5 is connected to B3 as max (0, 5-2) = 3
D4 is connected to B2 as max (0, 4-2) = 2
.
.
.
D0 is connected to B0 as max (0, 0-2) = 0

Similarly, the connections between instances of C and instances of D are determined as follows:

D5 is connected to C2 as max (0, 5-3) = 2 M

D0 is connected to C0 as max (0, 0-3) = 0

The process is shown in Figures 5(b) and 5(c).

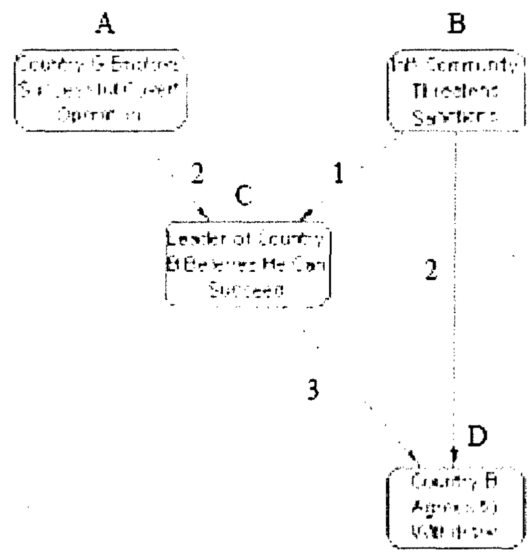

Figure 4: A Simple Timed Influence Net

<u>Step 3:</u> Let the input scenario is given. This step adds additional nodes and links in the TSBN based upon the selected course of actions provided as an input scenario. Suppose in the input scenario, action A is taken at time 2 while action B is taken at time 1. The maximum time stamp associated with the actionable nodes is 2, therefore this step adds two more slices to the TSBN and connects the parents and children as described in the previous step. Furthermore, the connections are added between the nodes representing actionable events as explained in Step 3(c) of the algorithm. For instance A0 and A1 are connected but since A is taken at time 2 therefore there is no connection between A1 and A2. The connections, however, are drawn between A2 and A3 and between A3 and A4, etc. Similarly, as B is 2 1 3 2 A D C B 12 taken at time 1 therefore connections are made between B1 and B2, B2 and B3, etc. while B0 is not connected to B1. The final TSBN is shown in Figure 5(d).

Once a TSBN is obtained from the corresponding TIN, the task of real time execution monitoring can be accomplished by entering observations in the model that arrive during the execution of the selected course of action. Suppose in the model of Figure 5(d), the evidence regarding variable D is received. The evidence states that D happens to be true at time 4. In the figure, all the indices of D equal to or greater than 4 are set to true. Thus D4, D5, D6, and D7 in Figure 5(d) become the evidence nodes. This new information revises the belief about the state of other non-actionable variables (non-root nodes) in the problem domain. For instance, this information would change the initial belief about C at time 1, 2, and onwards. If the time associated with the new information is greater than the number of slices drawn in the TSBN then more slices could be added to it. For example, if the new information says that D occurs at time 9 then the system modeler can add few

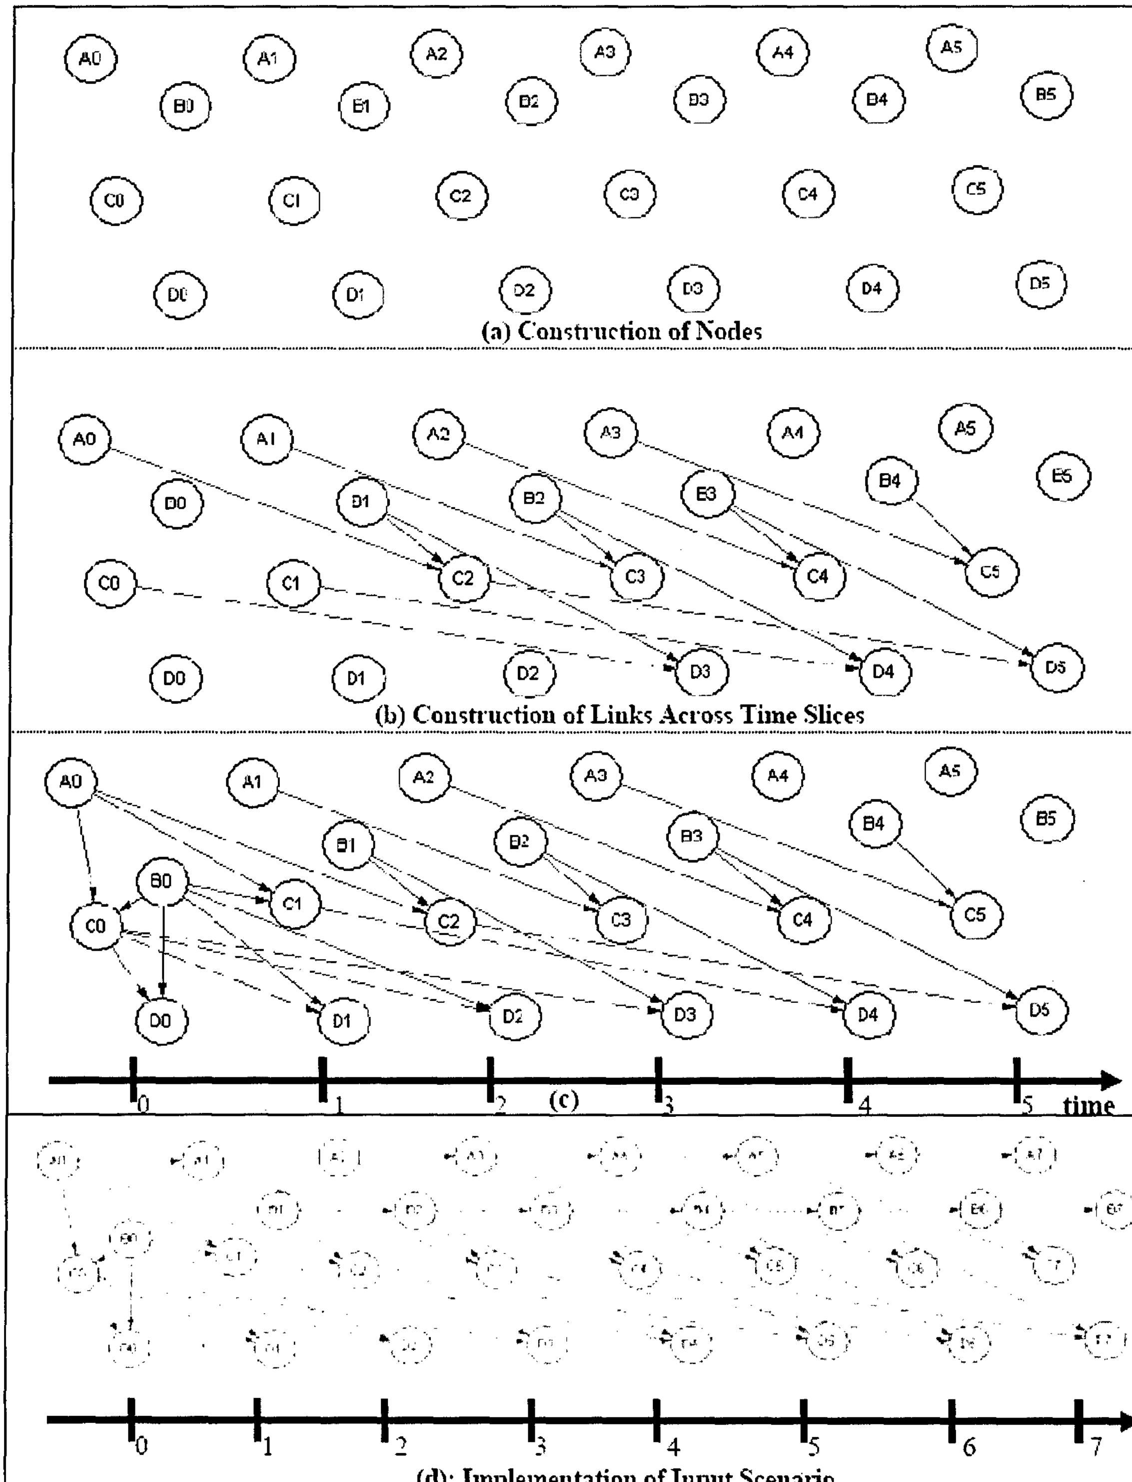

Figure 5: Steps Involved During the Transformation Process

more slices with the help of step 2 of the transformation algorithm in order to observe the impact of this new information on C at time 6 and onwards. It can be noticed in Figure 5(d) that variables in the TIN only depend upon the previous states of their parents and do not depend upon their own previous states. This is due to the fact that there is no link from node X_{t-1} to node X_t, where X belongs to the set of non-root nodes. There can be situations in which a variable's state at time t may depend upon its own state at a previous time stamp. In TSBNs, this issue is addressed by adding a link between different instances of the same variable at different time slices. The process is shown in Fig. 6 where the links in bold show the connection between (C_{t-1}, C_t) and (D_{t-1}, D_t). In TINs this requirement can be modeled by adding self-loops to such nodes (Fig. 6). This self-loop represents the dependency of the state of a variable at time t on its previous state.

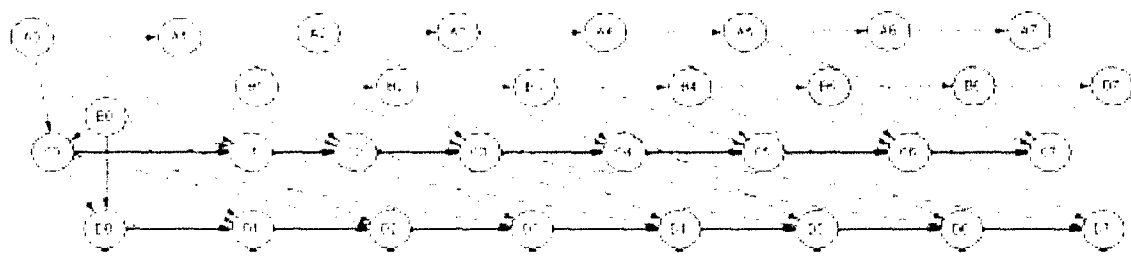

Figure 6: TSBN of Figure 5(d) with Dependencies Among Instances of Non-root Nodes

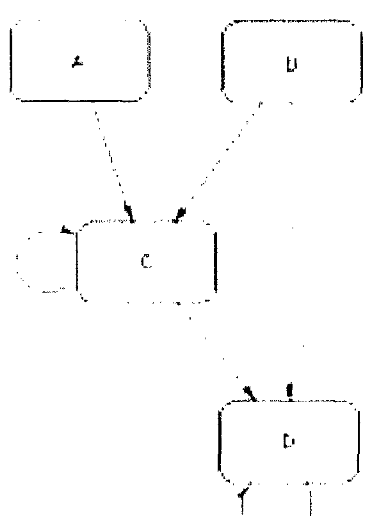

Figure 7: Timed Influence Net with Self-Loop

3.4 Conclusions

The paper presented a transformation algorithm for converting Timed Influence Nets into Time Sliced Bayesian Networks. The transformation provides the equivalence that exists between TINs and a class of TSBNs. Furthermore, the approach suggested in the paper delivers the advantages of both modeling paradigms to a system modeler. On one hand, it simplifies the knowledge elicitation process of TSBNs by suggesting TINs as a front end tool for modeling time and uncertainty; while on the other, it enhances the current capabilities of the TINs by providing the modeler ability to enter evidence that arrives during plan execution. In other words, the approach suggests that TINs be used for model building and course of action selection process, and TSBNs for execution monitoring of the selected course of actions. The task of inference in TSBNs, however, is computationally intractable. Thus, there is a tradeoff between the available approximate and exact algorithms in terms of accuracy and the time to compute probability of the variable of interest. The future research would focus on determining a set of inference algorithms (exact or approximate) that works better with the class of TSBNs that are obtained from TINs as a result of the transformation.

References

Agosta, J. M, "Conditional Inter-Causally Independent Node Distributions, a Property of Noisy- OR Models", In Proceedings of the 7th Conference on Uncertainty in Artificial Intelligence, 1991.

Boyen, X. and Koller, D. Tractable inference for complex stochastic processes. In Proc. of the Conf. on Uncertainty in AI, 1998.

Chang, K. C., Lehner, P. E., Levis, A. H., Zaidi, S. A. K., and Zhao, X., "On Causal Influence Logic", Technical Report, George Mason University, Center of Excellence for C3I, Dec. 1994.

Chariak, E., "Bayesian Network without Tears", AI Magzine, Winter 1991.

Doucet, A. de Freitas, N., Murphy, K. and Russell, S., "Rao-Blackwellised Particle Filtering for Dynamic Bayesian Networks", In Proceedings of the 16th Conference on Uncertainty in Artificial Intelligence, 2000.

Drudzel, M. J., and Henrion, M., "Intercausal Reasoning with Uninstantiated Ancestor Nodes", In the Proceedings of the 9th Conference on Uncertainty in Artificial Intelligence, 1993.

Figueroa, G. A., and Sucar, L. E., "A Temporal Bayesian Network for Diagnosis and Prediction", In Proceedings of the 15th Conference on Uncertainty in Artificial Intelligence, 1999.

Galan, S. F., and Diez, F. J., "Networks of Probabilistic Events in Discrete Time", International Journal of Approximate Reasoning, 30, pp. 181-202, 2002.

Haider, S. and Levis, A. H., "An Approximate Technique for Belief Revision in Timed Influence Nets", Proceedings of Command and Control Research and Technology Symposium, 2004.

Hanks, S., Madigan, D., and Gavrin, J., "Probabilistic Temporal Reasoning with Endogenous Change", In Proceedings of 11th Conference on Uncertainty in Artificial Intelligence, 1995.

Heckerman, D., and Breese, J. S., "Causal Independence for Probability Assessment and Inference Using Bayesian Networks", IEEE Transactions on Systems, Man, and Cybernetics – Part A: Systems and Humans, Vol. 26, No. 6, Nov. 1996.

Jensen, F. V., "Bayesian Networks and Decision Graphs", Springer-Veralg, 2001.
Kjaerulff, U., "A Computational Scheme for Reasoning in Dynamic Probabilistic Networks", In Proceedings of 8th Conference on Uncertainty in Artificial Intelligence, 1992.

Murphy, K and Weiss Y., "The Factored Frontier Algorithm for Approximate Inference in DBNs", In Proceedings of 17th Conference on Uncertainty in Artificial Intelligence, 2001.

Murphy, K., "Dynamic Bayesian Networks: Representation, Inference and Learning", PhD Thesis, UC Berkeley, Jul. 2002.

Neapolitan, R. E., "Learning Bayesian Networks", Prentice Hall, 2003.

Pearl, J., "Probabilistic Reasoning in Intelligent Systems: Network of Plausible Inference", Morgan Kaufmann, 1987.

Rosen, J. A., and Smith, W. L., "Influence Net Modeling with Causal Strengths: An Evolutionary Approach, Proceedings of the Command and Control Research and Technology Symposium, Naval Post Graduate School, Monterey CA, Jun. 1996.

Santos Jr. E., and Young, J. D., "Probabilistic Temporal Network: A Unified Frameword for Reasoning with Time and Uncertainty", International Journal of Approximate Reasoning, 2, 1999.

Takikawa, M., d'Ambrosio, B. and Wright, E., "Real-Time Inference with Large-Scale Temporal Bayes Nets", In Proceedings of the 18th Conference on Uncertainty in Artificial Intelligence, 2002.

SECTION 4

An Approximation Technique for Belief Revision in Timed Influence Nets

Sajjad Haider and Alexander H. Levis

4.1 Introduction

Probabilistic Belief Networks have gained popularity in last two decades to model uncertainty [Charniak, 1991], [Jensen, 2001], [Neapolitan, 2003], and [Pearl, 1987]. Commonly referred to as Bayesian Networks (BN), these belief networks use a graph-theoretic representation to explicitly show the dependencies among variables in a particular domain. Formally BNs are Directed Acyclic Graphs (DAG) in which nodes represent random variables while an edge connecting two nodes (typically) represents causal relationship (though it is not required that the connection be causal) between the two variables. The relationship between a node and its parents is defined by a Conditional Probability Table (CPT) for all combination of parents' states. The joint distribution over the random variables present in the network can be expressed as

$$P(x_1 \ldots\ldots x_n) = \prod_i P(x_i \mid pa(x_i))$$

where $pa(x_i)$ represents a configuration of the set of parents of variable x_i.

These networks have been primarily designed to simplify the intractable task of joint probability distribution elicitation. They have been usually applied without considering an explicit representation of time. In the past few years, efforts have been taken to integrate the notion of time and uncertainty [Figueroa and Sucar, 1999], [Galan and Diez, 2002], [Hanks et al., 1995], [Kjaerulff, 1992], and [Santos and Young, 1999]. The popular approach of modeling processes having temporal dependencies is to discretize the time and create an instance of each random variable present for each point in time. The process starts with eliciting the probability distributions for the static probabilistic model. This model is repeated for difference time slices, and links are drawn between these slices to represent the temporal dependencies among the nodes in the network. The approach is usually referred to as Time Sliced Bayesian Networks (TSBNs) or Dynamic Bayesian Networks (DBNs) [Murphy, 2002]. Figure 1 shows two types of TSBNs as discussed by Hanks et al., [1995]. In Figure 1 on the left, all the connections in the models are inter-slice, i.e., connections only exist among variables within different time slices. On the contrary in Figure 1 on the right, the variables in the model have both inter-slice and intra-slice connections.

Figure 1: Examples of Time Sliced Probabilistic Networks

Despite the fact that probabilistic belief networks address the intractable problem of eliciting joint distribution of random variables in an efficient way, the number of parameters required to specify the Conditional Probability Table (CPT) of a node increases exponentially with the number of parents. Several approaches have been proposed that estimate the CPT values from parameters that are linear with the number of parents. These include but not limited to Noisy-OR [Agosta, 1991], [Drudzel and Henrion, 1993], and [Heckerman and Breese, 1996], CAusal STrength (CAST) Logic [Chang et al., 1994] and [Rosen and Smith, 1996], etc. The probabilistic models that use CAST Logic as an interface for estimating the CPTs for each node in the network are referred to as Influence Nets. Influence Nets simplify knowledge elicitation by reducing the number of parameters that must be provided. They are appropriate for modeling situations in which the estimate of the conditional probability is subjective, e.g., when modeling potential human reactions and beliefs, and when subject matter experts find it difficult to fully specify all conditional probability values.

Timed Influence Nets (TINs) [Wagenhals and Levis, 1999], extend the CAST logic based interface of Influence Nets by providing a way to model uncertainty and temporal constrains present in a stochastic model from a Discrete Event System's (DES) perspective. These TINs are developed by making cause and effect or influencing relationships among variables in the domain. The links between two variables represents the temporal causal relationship between them. The impact of one variable on other variables does not necessarily occur instantaneously; rather it may occur after a specified time. This time is represented by the assignment of a delay to each link. All the nodes in the network could also have an optional time delay which represents the information processing delay of the corresponding node. The marginal probability of a node is computed whenever there is a change in the state (the marginal probability) of any of its parents. To achieve this behavior in a computationally efficient manner, TINs propagate probabilities using independence of parents assumption, also referred to as loopy belief propagation [Murphy et al., 1999].

TINs have been used experimentally in the area of Effect Based Operations (EBO) [Wagenhals et al., 2003]. They are modeled by identifying target variables and relating them to the actions that could impact them. The purpose of creating such models is to determine how to maximize or minimize the probability of occurrence of the target variables by taking a timed sequence of actions or actionable events. Actionable events, in this context refer to the random variables that are modeled as root nodes in the corresponding TIN. The actionable events (either under the control of the decision maker or the adversary) and the variable of interest (target variable) are connected through chains of variables that represent intermediate effects. Some of these variables may be observable. This paper describes an extension to the capability of TINs by adding the provision of incorporating new evidence in the model. The algorithm presented in the paper provides an approximation scheme for updating the belief of the affected variables after observing evidence provided that certain constraints are satisfied. The algorithm first tries to reduce the net by identifying those variables which are relevant for computing the posterior probability of a target variable over an interval of time. The nodes, which do not have any impact on the variable of interest as a result of new evidence, are considered as being pruned from the net. In the second step, the algorithm computes the new beliefs on all the affected variables taking into account the time delays (communication and processing delay) present in the graph. This technique provides an initial step in the direction of integrating the impact of planned Course of Action (COA) [Wagenhals and Levis, 2000] selected by the decision maker over a time interval and the information/observations which arrive during / after the execution of that particular COA. The objective is to assess impact of the actionable events as the situation dynamically unfolds.

The rest of the section is organized as follows. Section 4.2 discusses the mathematical formulation of TINs and their application. Section 4.3 describes the belief propagation algorithm that supports the incorporation of evidence, while Section 4.4 concludes the paper and points out areas for future research.

4.2 Modeling Uncertainty Using Timed Influence Nets

4.2.1 Knowledge Elicitation

TINs use CAST logic, a variant of Noisy-OR, to simplify knowledge elicitation from subject matter experts. Instead of assigning conditional probabilities, the expert first specifies the qualitative relationship between two connected nodes as either promoting or inhibiting. Figure 2 shows a simple two node influence net. In Figure 2(a), the modeler indicates that the presence of A can cause B (with some probability), and a '+' sign is attached to the arc. Similarly, the modeler indicates that the absence of A can inhibit B, therefore there is a '-' sign attached in Figure 2(b). Figure 2(c) shows the aggregated qualitative relationship between two nodes by using the double (+, -) notation. If the modeler had determined that the presence of A inhibits B while the absence of A promotes B, then the aggregate notation would be (-, +). Qualitative relationships among variables have also been applied for Qualitative Probabilistic Networks (QPN) [Drudzel and Henrion, 1993] and Causal Maps [Nadkarni and Shenoy, 2001].

After (or during) assigning the qualitative relationships between the two nodes, the expert(s) quantify these relationships by assigning values on the scale of 0 to 1. Low values mean the promoting or inhibiting relationship is weak while values near 1 mean the relationship is strong. The CAST logic uses a heuristic to convert theses qualitative relationships to form the conditional probability matrix for each non-root node. Besides reducing the number of parameters required for specifying the conditional probability matrix for each node, the CAST logic also helps in eliciting knowledge from different subject matter experts. For instance, in an international conflict, there are many dimensions of the problem, namely political, religious, ethnic, social, etc. It is difficult to find domain experts having expertise in all the above areas. The CAST logic provides a mechanism to obtain information from various experts and then combine their individual assessment in a mathematical manner. The exact details of the CAST logic algorithm are beyond the scope of this paper. The interested reader should refer to Chang et al., [1994] and Rosen and Smith, [1996].

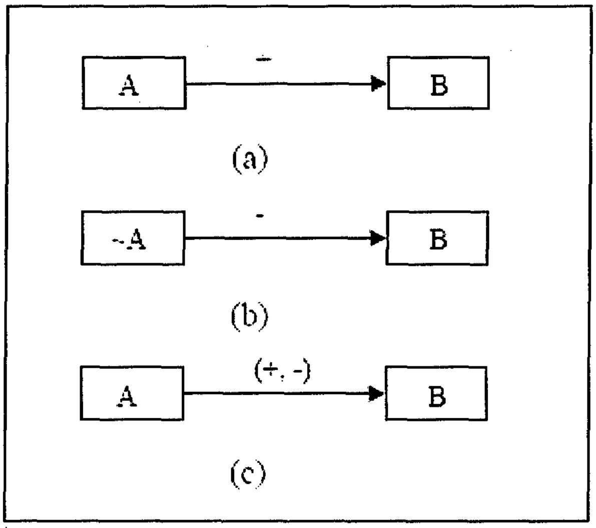

Figure 2: Qualitative Relationships in TIN

Timed Influence Nets extend the capabilities of Influence Nets by providing a mechanism to specify certain kinds of temporal constraints present in a problem domain. Wagenhals et al. [2003] have classified 4 types of temporal information that could be associated with a Timed Influence Nets. Out of them, three are part of the model itself and one is related to the input scenario. The *input scenario* can be described in terms of the actions in the Course of Action (COA) and the time at which these actions occur. Among the remaining three types of temporal information, one is related to the *duration* of an action. The second type is related to the *communication* and *processing* delay present in a problem domain. In other words, this type represents the amount of time it takes for knowledge about a change in the status of any variable to be propagated by some real world phenomenon to the node that is affected by that change. The third type of temporal information is sometimes referred to as *persistence*. This is the time interval over which an effect is manifested. Because of the complexity of this problem, the issue of modeling persistence is still an area of active research. In the sequel, when we discuss TINs (a) (b) (c) A B + ~A B - A B (+, -) we mean Influence Nets that are capable of modeling the first three types of temporal information (without persistence). The full specification of a Timed Influence Net is as follows

1. A set of random variables that makes up the nodes of a TIN. All the variables in the TIN have binary states.
2. A set of directed links that connect pairs of nodes.
3. Each link has associated with it a pair of CAST Logic parameters that shows the causal strength of the link (usually denoted as g and h values).
4. Each non-root node has an associated CAST Logic parameter (denoted as baseline probability) while each root node has a prior probability.
5. Each link has a corresponding delay d (where d > 0) that represents the communication delay.
6. Each node has a corresponding delay e (where e > 0) that represents the information processing delay.
7. A pair (p, t) for each root node, where p is a list of real numbers representing probability values. For each probability value, a corresponding time interval is defined in t. In general, (p, t) is defined as

$$([p_1, p_2, \ldots, p_n], [[t_{11}, t_{12}], [t_{21}, t_{22}], \ldots, [t_{n1}, t_{n2}]])$$

$$\text{where } t_{i1} < t_{i2} \text{ and } t_{ij} > 0 \qquad \forall i = 1, 2, \ldots, n \text{ and } j = 1, 2$$

The first four requirements in the above specifications are the same for static and timed Influence Nets. The last three requirements are related specifically to TINs. Once a TIN is completely specified, it can be used to observe the behavior of variables of interest over a specific period of time.

4.2.2 Course of Action Analysis

Figure 3 [Wagenhals and Levis, 1999] shows how a TIN model compactly represents actionable events, causal or influencing relationships between actions and effects, the strengths of those relationships, and the time delays associated with effect propagation. It illustrates the kind of analysis that could be done using TINs. The model shows the cause and effect relationship as seen by an expert on international politics.

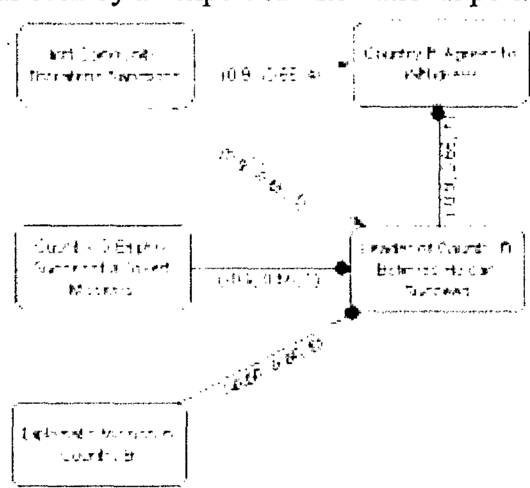

Figure 3: A Simple Timed Influence Nets

Country B has occupied portion of land of a neighboring country. The objective is to find a peaceful solution of the problem, or, in other words, the objective is to determine the probability that country B would agree to withdraw its forces. There are five variables in the Influence Net represented by the five boxes. Each arc in the net has an annotation that is a triple. The first two elements of the triple contain the influence strength of the presence and absence of the parent node on the child node. The third entry is the time delay required for the influence to reach from parent to child. The prior probabilities of each root node are also shown in the figure.

The next stage is to set the time for execution of the actions represented as root nodes. Suppose the actions are executed at the following times:

Diplomatic Mission in Country B @ 8
Int'l Community Threatens Sanctions @ 10
Country G Employs Successful Covert Mission @ 11

The influences of these actions reach the target node (Country B Agrees to Withdraw) at different times. Every time an influence arrives at the target node, the TIN updates the belief of target node. These belief updates and the time each one occurs are shown in the form of a probability profile (Figure 4).

4.3 Belief Propagation Algorithm

TINs were originally designed for the COA analysis. In the original TIN formulation it was assumed that all the actions are in fact evidence nodes and there would be no evidence on the other nodes in the networks. Thus, these TINs lacked the ability to incorporate the information/evidence coming from different sources during the execution of a COA. In a military/political scenario, this new information might come from the surveillance system observing an adversary's actions. In an economic domain, a new development in the market, e.g., bankruptcy filed by some corporation, might be taken into account before making a strategic decision. In any case, this new information results in the revision of a previously held belief about some variables in the domain. This section, which is the main theme of this paper, extends the ability of the original TINs by presenting an approach for integrating the new information with the existing beliefs on other variables in the net. The algorithm takes advantage of the fact that while analyzing a TIN, the analyst is primarily interested in observing the behavior of the desired objectives. This feature helps in applying graph reduction techniques and simplifies the belief revision process. Instead of revising the beliefs on all the variables in the net, the belief revision process is applied to only those set of variables which impact the variable of interest in some way. The algorithm is based on the constraint *that the marginal probability of a parent node will not be updated unless all of its children which need to be updated have updated marginal probabilities.*

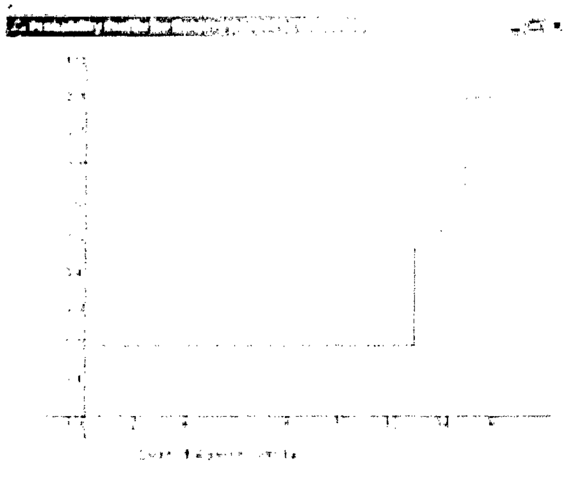

Figure 4: Probability Profile for Node "Country B Agrees to Withdraw"

The algorithm has three main steps. The first step determines the sequence in which the marginal probability of nodes will be updated once evidence has been incorporated in one of the nodes of the model. Step two selects only the nodes in the sequence that are needed to update the target node. In Step 3, the updating is accomplished node by node in the sequence determined in Step one. The following sub sections describe the working of the algorithm in detail.

4.3.1 Sequencing of Nodes

The algorithm that determines the sequence for updating the nodes in the TIN is presented in Table 1. It creates a sequenced list of all of the ancestors of a node to which evidence is applied. The sequence is based on a breadth first protocol that ensures the closest ancestors are placed on the list before more distant ancestors. This sequence shows the order of the belief updating of nodes in the net. The node for which we have obtained hard evidence is assigned sequence number one. The parents of the node are sequenced then, and the process continues until all root nodes that have a path to the evidence node are reached. We call this process backward sequencing. The resultant sequence is then used to update the probabilities of nodes during backward propagation. It should be noted that after reaching a root node, the belief updating process continues in the forward direction and the nodes, which were not updated during the backward direction, are then updated until the algorithm reaches the target node. The example TIN in Figure 5(a) is used to explain the sequencing algorithm. Only the structure of the TIN and the time delays are shown in the figure for clarity.

Suppose evidence about node 'H' in Figure 5(a) is obtained. In the first step, the algorithm initialized the NodeList with 'H'. In the second step, the variable 'Current_Node' is assigned the first unprocessed node in 'NodeList', which in this case is 'H'. Since 'H' is the evidence node, the parents of the 'H' are entered in the list, and 'H' is considered to be a processed node. At the end of the first iteration of step 2, 'NodeList' has value [H, I, F].

Table 1: Algorithm for Sequencing of Nodes

```
Let X be the evidence node.
1) NodeList = [X]
2) While NodeList has unprocessed nodes
      Current_Node  =  1st Unprocessed node in the
                            NodeList
   If there exists descendant of Current_Node
   unprocessed in the NodeList then
      Move Current_Node at the end of the NodeList
   Else
      Add parents of Current_Node in the NodeList
      Mark Current_Node as processed
 End Loop
```

Nodes 'I' and 'F' are unprocessed. The sequence of choice of parents of 'H' is arbitrary, so 'NodeList' could have the value [H, F, I]. In the next iteration, 'I' becomes the Current_Node. None of the descendants of 'I', are in the list of unprocessed nodes, so the parent of 'I' (node 'D') is added to in the 'NodeList'. At the end of this iteration 'NodeList' has the value [H, I, F, D]. The next unprocessed variable in the list is 'F' and since the descendants of 'F' have already been processed in the list, the parents of 'F' are included in the list. The 'NodeList' now has value [H, I, F, D, G, B]. Next 'D' is assigned to Current_Node but one of the descendants of 'D', i.e., 'G', is still in the list unprocessed. Thus 'D' is moved to the end of the list. The new value of 'NodeList' is [H, I, F, G, B, D]. This time node 'G' is the Current node and the only child of 'G' has already been included in the list, therefore parents of 'G' are also included in the list, making the value of NodeList be [H, I, F, G, B, D, E]. In this way, the algorithm keeps iterating, until all the nodes in the NodeList are processed. The final value of NodeList is [H, I, F, G, E, D, C, B, A, M]. In general, there can be more than one possible sequence, however, all sequences will produce the same results as far as belief updating is concerned.

4.3.2 Graph Reduction

The steps described in the previous sections give the node ordering that would be used while updating the nodes in the backward direction. But not all of the nodes are required, if the objective is to only see the impact of the evidence on the target node. Considering the same example used in the previous section, nodes 'A' and 'M' represent the actions taken by the decision maker(s). If the evidence is received after the execution of these actions and the objective is to analyze the behavior of node 'K', then there is no need to update nodes G, E, D, and C. All we need to do is to update I, F, and B in the backward propagation and then update the descendants of these variables during forward propagation. The process is referred as graph pruning. The resulting TIN is shown in Figure 5(b). The sequencing and pruning algorithms can be used when evidence is available for more than one node. When evidence is available at two or more nodes then the sequence and pruning algorithms are run multiple times. For example, if there is evidence for nodes H and L, then only nodes I and J need to be updated in order to see the effect of the evidence of both nodes on node 'K'. The remaining nodes do not need to be updated, as they could not impact 'K' through some other paths.

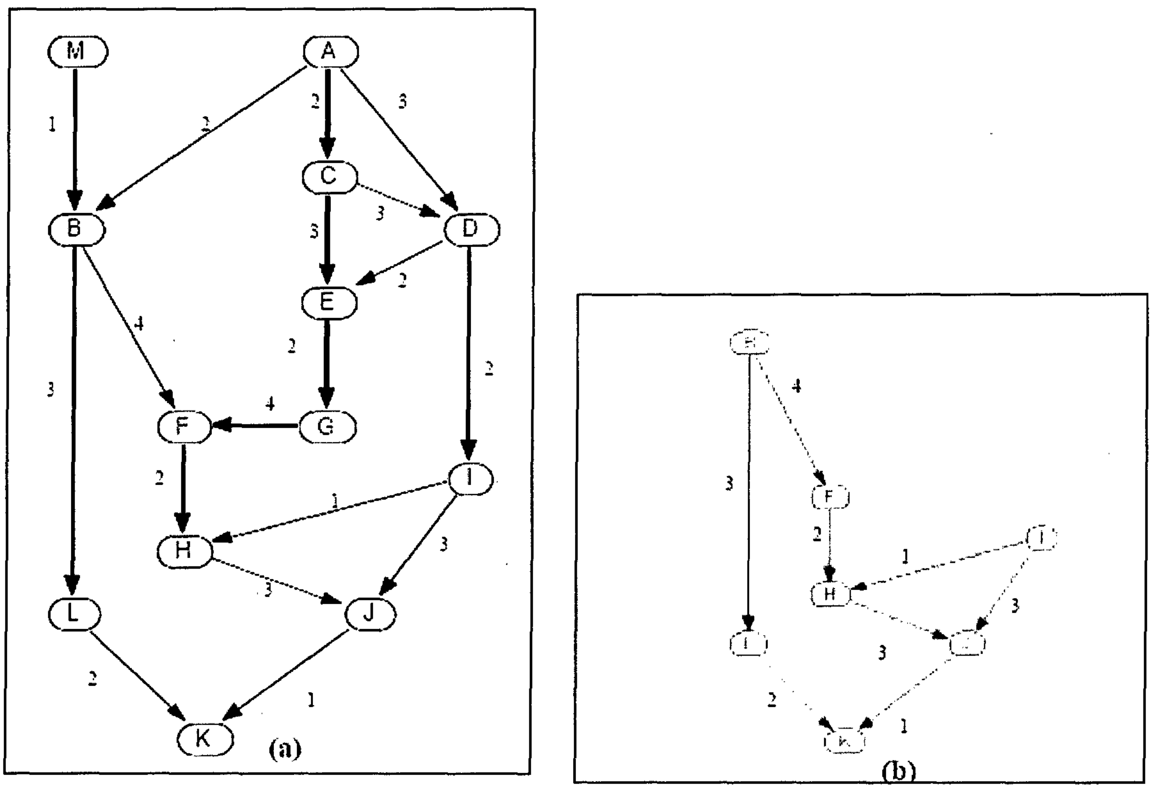

Figure 5(a): TIN for Explaining the Sequencing of Nodes Algorithm
Figure 5(b): TIN after Pruning

4.3.3 Computation of Posterior Probability

Once the sequence is obtained, the iterative application of Bayes' rule is used for computing the posterior distribution of the affected variables in the net. In our example, suppose the decision maker has taken actions 'M' and 'A' at time 6 and 8, respectively. The conditional probability tables associated with each non-root node are not shown in the figure to enhance the readability of the figure. The reader should be able to trace the flow of information as the actions take place. For instance, node 'B' would be updated at time 7 and 9. Similarly, node 'D' would be updated at times 11 and 13. The times of update for other nodes could be computed in the similar manner. Once we finish the update in the forward direction, the findings could be entered in the system. Suppose we get evidence about the presence of 'H' at time 24. Mathematically, P'(H) = 1.0 @ 24, where the notation P'(H) means that this is an updated marginal probability. Before getting this evidence, the marginal probabilities and their times of update for node 'H' and its parents 'F' and 'I' are as follows:

$P(H) = 0.44$ @ 23
$P(F) = 0.21$ @ 21
$P(I) = 0.12$ @ 15

After updating node 'H' at time 24, the belief updating process selects the next element in 'NodeList' which is 'I' in our example. The time delay on the arc between 'H' and 'I' is 1. Hence the probability of 'I' at time 23 should be revised. The probability is computed as

$$P'(I) = P(I \mid H)\, P'(H) + P(I \mid \sim H)\, P'(\sim H) \qquad (1)$$

where $P(I \mid H)$ and $P(I \mid \sim H)$ are computed using Bayes' rule. $P'(H)$ and $P'(\sim H)$ represent the new probability of H and ~H as a result of the new evidence.

$$P(I \mid H) = P(H, I) / P(H) \qquad (2)$$

$$P(I \mid \sim H) = P(\sim H, I) / P(\sim H) \qquad (3)$$

The numerator of Eq. (2) can be expanded as

$$P(H, I) = P(H, I, F) + P(H, I, \sim F) = P(H \mid I, F)\, P(I, F) + P(H \mid I, \sim F)\, P(I, \sim F) \qquad (4)$$

As discussed earlier, TINs assume that both parents of 'H', i.e., 'I' and 'F', are independent, which results in the simplification of equation (4),

$$P(H, I) = P(H \mid I,F)\, P(I)\, P(F) + P(H \mid I,\sim F)\, P(I)\, P(\sim F) \qquad (5)$$

Suppose the Conditional Probability Matrix for node 'H' is given as

$P(H \mid I, F) = 0.15$ $\quad$ $P(H \mid I, \sim F) = 0.98$

$P(H \mid \sim I, F) = 0.005$ $\quad$ $P(H \mid \sim I, \sim F) = 0.5$

Using these conditional probabilities, equation (5) becomes

$$P(H,I) = (0.15)(0.12)(0.21) + (0.98)(0.12)(0.79) = 0.097$$

The numerator of equation (3) can be computed in a similar manner

$$P(\sim H,I) = (0.85)(0.12)(0.21) + (0.02)(0.12)(0.79) = 0.023$$

From the above two equations, we can compute equations (2) and (3)

$$P(I \mid H) = 0.097 / 0.44 = 0.22$$

$$P(I \mid \sim H) = 0.023 / 0.56 = 0.04$$

Equation (1) thus becomes:

P'(I) = (0.22)(1.0) + (0.04)(0) = 0.22

The probability of F is updated in a similar manner at time 22, as the time delay on the arc between 'F' and 'H' is 2. After these updates, the probabilities of H, F, and I become:

P'(H) = 1.0 @ 24

P'(F) = 0.01 @ 22

P'(I) = 0.22 @ 23

The next node in the 'NodeList' is 'B'. The time delay between 'F' and 'B' is 4. Hence the probability of B is revised at time 18. It has been discussed earlier that there is no need to update the probability of nodes 'G', 'E', 'D', and 'C' if the only objective is to observe the impact of evidence on the target node 'K'. But in order to show how the constraint of not updating the parent unless all the children are being updated works, we could continue the process of backward propagation till we reach node 'D'. The impact of evidence arrives at 'D' through both of its children, 'I' and 'E' at time 21 and 14, respectively. If we update 'D' at 21 based on the new probability of 'I' at time 23 then the probability of 'E' will also be updated at 23 during forward propagation. Further, the probability of 'F' would be updated at time 29. Since the child of 'F', i.e., 'H' is already an evidence node, the new probability at 'F' results in the update of the other parent of 'H', namely 'I' at time 30. Thus, as a result of updating the probability of 'D' based on the new value of 'I' we have obtained a new probability of 'I'. This cycle of update would continue forever. In order to avoid falling into the problem of infinite loop which would result if we consider the impacts of 'I' and 'E' on node 'D' individually, the earliest time will be used for the update. Hence node 'D' will be updated at time 14. In general, if the impact of evidence reaches node Y through multiple paths then the update time is computed as

$$t_Y = \min(t_{X1} - \alpha_{YX1}, t_{X2} - \alpha_{YX2}, \ldots\ldots, t_{XN} - \alpha_{YXN})$$

where *X1, X2, ……, XN* represent the children of Y that are already updated. α_{YX1}, α_{YX2}, … …, α_{YXN} represent the time delay on the links between Y and its children *X1, X2, … … XN*, respectively while t_{X1}, t_{X2}, …., t_{XN} represent the time of update of *X1, X2, … … XN*, respectively.

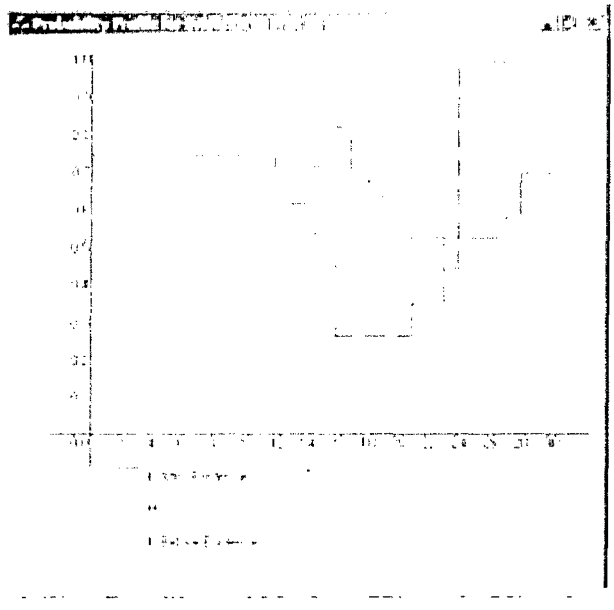

Figure 6: Probability Profiles of Nodes 'H' and 'K' After Evidence on 'H'

Once the backward propagation is finished, the algorithm starts in the forward direction. The probability of node 'L' is updated at time 21 (the time of last update of B plus the time delay on the arc between 'B' and 'L'). The impact of evidence on node 'J' arrives through two paths: H-I-J, and H-J. The last update of I occurred at time 23 and the arc delay has value 3. Hence there is a change in the marginal probability of 'J' at time 26. The impact of evidence at 'H' reaches 'J', through the direct path between 'H' and 'J', in 3 time units. Thus, there is another update at time 27 (time of evidence plus delay on arc from H to J). The changes in the marginal probabilities of the parents of the target node 'K' would result in the computation of new beliefs in node 'K' at time 23, 27, and 28. These changes are shown in the probability profile of Figure 6. The figure shows that incorporating the evidence about node H at time 24 changes the probability of node K from 0.57 to 0.70 at time 28. This example assumes that the time delays are associated only with the arc. The same technique can be applied for computing the posterior probability if the time delays are associated with both nodes and arcs.

The ability to enter evidence in the model allows the possibility to compute the value of information. In the planning domain, the decision makers are confronted with the task of the placement of evidence gathering resources that may be limited in number. Having the ability of finding the impact of certain evidence on the desired objective, the planners would be in a better position to decide how to use these scarce resources based on the contribution each evidence node makes in reducing the uncertainty in the objective node.

4.4 Conclusions

This section presented a computationally efficient technique for belief updating in Timed Influence Nets. The proposed technique updates the nodes in the sequential manner using the constraint that all the children of a node affected by the new evidence will be updated first before updating the belief in that particular parent node. The algorithm also takes advantage of the fact that in TINs, the focus is on observing the behavior of few nodes in the network. Hence there is no need to update all the nodes of the network. The nodes that receive impact of evidence and have a path to the target nodes only need to be updated. This relaxation helps in applying graph reduction techniques on TINs.

One of the possible limitations of the approach is that the algorithm works only if the time stamp of the evidence is later than the time stamp of the last update of the evidence node caused by the forward propagation of the effects from all of the action nodes. This constraint might be very strong in some cases. It is quite possible that the evidence could be observed earlier than expected by the model. There are few possible approaches to relax this constraint. Either the expert should revise the communication and processing delays in the network, or the portion of the graph which is in conflict with this new evidence should be pruned before starting the backward propagation. The other alternative is to convert the Timed Influence Nets into a Time Sliced Bayesian Network (TSBN) and use a vast variety of algorithm available for TSBN. The transformation from TINs to TSBNs is addressed in a future paper. It should be mentioned, though, that the problem of inference in TSBNs (or even in static Bayesian Networks) is computationally intractable. Thus, there is a trade off between the use of approximate algorithms and the exact algorithms in terms of accuracy and the time to compute the posterior probability of the variable of interest. An efficient algorithm for belief updating in TSBNs is still an area of active research.

References

Agosta, J. M, "Conditional Inter-Causally Independent Node Distributions, a Property of Noisy-OR Models", In Proceedings of the 7th Conference on Uncertainty in Artificial Intelligence, 1991.

Chang, K. C., Lehner, P. E., Levis, A. H., Zaidi, S. A. K., and Zhao, X., "On Causal Influence Logic", Technical Report, George Mason University, Center of Excellence for C3I, Dec. 1994.

Chariak, E., "Bayesian Network without Tears", AI Magzine, Winter 1991.

Drudzel, M. J., and Henrion, M., "Intercausal Reasoning with Uninstantiated Ancestor Nodes", In the Proceedings of the 9th Conference on Uncertainty in Artificial Intelligence, 1993.

Drudzel, M. J., and Henrion, M., "Belief Propagation in Qualitative Probabilistic Networks", In N. Piera Carrete and M. G. Singh (Eds.), Qualitative Reasoning and Decision Technologies, pp. 451 – 460, CIMNE, Barcelona, 1993.

Figueroa, G. A., and Sucar, L. E., "A Temporal Bayesian Network for Diagnosis and Prediction", In Proceedings of the 15th Conference on Uncertainty in Artificial Intelligence, 1999.

Galan, S. F., and Diez, F. J., "Networks of Probabilistic Events in Discrete Time", International Journal of Approximate Reasoning, 30, pp. 181-202, 2002.

Hanks, S., Madigan, D., and Gavrin, J., "Probabilistic Temporal Reasoning with Endogenous Change", In Proceedings of 11th Conference on Uncertainty in Artificial Intelligence, 1995.

Heckerman, D., and Breese, J. S., "Causal Independence for Probability Assessment and Inference Using Bayesian Networks", IEEE Transactions on Systems, Man, and Cybernatics – Part A: Systems and Humans, Vol. 26, No. 6, Nov. 1996.

Jensen, F. V., "Bayesian Networks and Decision Graphs", Springer-Veralg, 2001.

Kjaerulff, U., "A Computational Scheme for Reasoning in Dynamic Probabilistic Networks", In Proceedings of 8th Conference on Uncertainty in Artificial Intelligence, 1992.

Murphy, K., Weiss, Y., and Jordan, M., "Loopy-belief Porpagation for Approximate Inference: An Empirical Study", In Proceedings of the 15th Conference on Uncertainty in Artificial Intelligence, 1999.

Murphy, K., "Dynamic Bayesian Networks: Representation, Inference and Learning", PhD Thesis, UC Berkeley, Jul. 2002.

Neapolitan, R. E., "Learning Bayesian Networks", Prentice Hall, 2003.

Pearl, J., "Probabilistic Reasoning in Intelligent Systems: Network of Plausible Inference", Morgan Kaufmann, 1987.

Rosen, J. A., and Smith, W. L., "Influence Net Modeling with Causal Strengths: An Evolutionary Approach, Proceedings of the Command and Control Research and Technology Symposium, Naval Post Graduate School, Monterey CA, Jun. 1996.

Santos Jr. E., and Young, J. D., "Probabilistic Temporal Network: A Unified Frameword for Reasoning with Time and Uncertainty", International Journal of Approximate Reasoning, 2, 1999.

Nadkarni, S. and Shenoy, P. P., "A Bayesian Network Approach for Making Inference in Causal Maps", European Journal of Operation Research, vol. 128, No. 3, 2001.

Wagenhals, L. W., and Levis, A. H., "Converting Influence Nets with Timing Information to a Discrete Event System Model, A Colored Petri Net", Proc. Of 2nd Workshop on Practical Uses of Colored Petri Nets and Design/CPN, DAIMI PB-532, Aarhus University, Denmark, 1999.

Wagenhals, L. W. and Levis, "Course of Action Development and Evaluation, Proceedings of Command and Control Research and Technology Symposium, 2000.

Wagenhals, L. W., Levis, A. H., and McCrabb, M. B., "Effects-Based Operations; A Historical Perspective of a Way Ahead", Proceedings of 8th International Command and Control Research and Technology Symposium, 2003.

SECTION 5

Dynamic Influence Nets: An Extension of Timed Influence Nets for Modeling Dynamic Uncertain Situation

Sajjad Haider and Alexander H. Levis

5.1 Introduction

Decision making in uncertain complex situations has always been a very difficult task. Access to a large amount of information has further magnified the complexity of this problem. For an *organization*, it has become an unmanageable task to analyze enormous amounts of information in a timely manner. Several efforts have been made to model this problem using the framework of probabilistic reasoning and inferencing, commonly referred to as Bayesian Networks [Pearl, 1987]. Timed Influence Nets (TINs), one of the instances of this framework, have been used experimentally in the area of Effects-Based Operations (EBOs). [Wagenhals and Levis, 2002;Wagenhals et al., 2003;Wagenhals and Wentz, 2003] A TIN models the uncertainties and temporal constraints present in a stochastic domain from a Discrete Event System (DES) perspective. The purpose of building a TIN is to form a coherent model of the situation at hand by combining the knowledge of several subject matter experts. The resultant model is then used as an aid to decision makers to help them reach a final *decision* in a rational manner.

Despite their acceptance as a modeling tool, the assumptions made by TINs limit the capabilities of a system modeler in terms of expressing real world situations. For instance, TINs do not model the impact of the sequence in which actions are taken. Thus, no matter what the sequence of the actions is, the final probability of achieving the desired *effect(s)* remains same. This behavior is because of the underlying assumption in TINs that events are *memoryless*, i.e. the probability of occurrence of an event at a particular time instant does not depend upon its own probabilities of occurrence during the previous time instants. As a consequence, the probability of an event depends only upon the actions executed so far and not on the sequence in which these actions are executed. This paper is an attempt to overcome this limitation of TIN.

Moreover, the assumption made by TINs that the influence of a cause remains the same once the cause has occurred is found to be unrealistic. In reality, events happen and they influence other relevant events. In many cases, the intensity of their influences decays over time. Thus, an event having a very strong influence at the time of its occurrence on another event might have an insignificant influence after a certain period of time. For example, a resolution passed by the United Nations has a very strong impact on the concerned parties at the time of its approval. As the days pass, the resolution starts loosing its affect and after a couple of months it completely looses its importance unless the problem is solved or it is backed up by another resolution on the same subject. That is, the influence of an event is *time-varying*. Currently, TIN lacks the ability to model such cases. It assumes that the causal strength of the influences does not change over time. This paper extends the capabilities of TIN by proposing a way for modeling time-varying influences.

The rest of the paper is organized as follows. Sections 2 and 3 describe Influence Nets and Timed Influence Nets, respectively. Sections 4 and 5 describe the limitations of TINs and the proposed enhancement. Finally, Section 6 concludes the paper and points towards the future research direction.

5.2 Influence Nets

Influence Nets (INs), an instance of the Bayesian framework, were proposed a decade ago to overcome the intractability issues present in BNs. They employ an approximation inference algorithm, termed as loopy belief propagation [Kschischang and Frey, 1998;McEliece et al., 1998;Murphy et al., 1999], and non-probabilistic knowledge acquisition interface, termed as the CAST logic [Chang et al., 1994;Rosen and Smith, 1996].

The modeling of the causal relationships using an IN is accomplished by connecting a set of actionable events and a set of desired effects through chains of cause and effect relationships. The strength of such relationships is specified using the CAST logic parameters (a brief overview of the logic is presented later in this section) instead of the probabilities. The required probabilities are internally generated by the CAST logic with the help of user-defined parameters. The Influence Nets are therefore appropriate for the following situations: i) for modeling situations in which it is difficult to fully specify all conditional probability values ii) and/or the estimates of conditional probabilities are subjective, and iii) estimates for the conditional probabilities cannot be obtained from empirical data, e.g., when modeling potential human reactions and beliefs.

The actionable events in an IN are drawn as root nodes (nodes without incoming edges). A desired effect, or an objective the decision maker is interested in, is modeled as a leaf node (node without outgoing edges). Typically, the root nodes are drawn as rectangles, while the non-root nodes are drawn as rounded rectangles. Consider the IN of Figure 1. The text associated with the non-root nodes represents the corresponding conditional probability values obtained from the CAST logic parameters (not shown in the figure) while the text associated with the root nodes represents the prior probabilities. The texts associated with arcs are time delays and are explained in Section 3. The belief propagation scheme used in INs is based on independence of parents assumptions. Thus, the marginal probability of a non-root node is computed with the help of its Conditional Probability Table (CPT) and the prior probabilities of its parents. For instance, the marginal probability of variable A is computed as

$$P(A) = P(A \mid \neg B, \neg E)P(\neg B)P(\neg E) + P(A \mid \neg B, E)P(\neg B)P(E) + P(A \mid B, \neg E)P(B)P(\neg E) + P(A \mid B, E)P(B)P(E)$$
$$= 0.005 \times 0.95 \times 0.99 + 0.95 \times 0.95 \times 0.01 + 0.95 \times 0.05 \times 0.99 + 0.99 \times 0.05 \times 0.01$$
$$= 0.06$$

The probability of D is then computed by using its CPT and the marginal probabilities of A (computed above) and E. Thus

$$P(D) = P(D \mid \neg E, \neg A)P(\neg E)P(\neg A) + P(D \mid \neg E, A)P(\neg E)P(A) + P(D \mid E, \neg A)P(E)P(\neg A) + P(D \mid E, A)P(E)P(A)$$
$$= 0.05 \times 0.99 \times 0.94 + 0.95 \times 0.99 \times 0.06 + 0.001 \times 0.01 \times 0.94 + 0.05 \times 0.01 \times 0.06$$
$$= 0.11$$

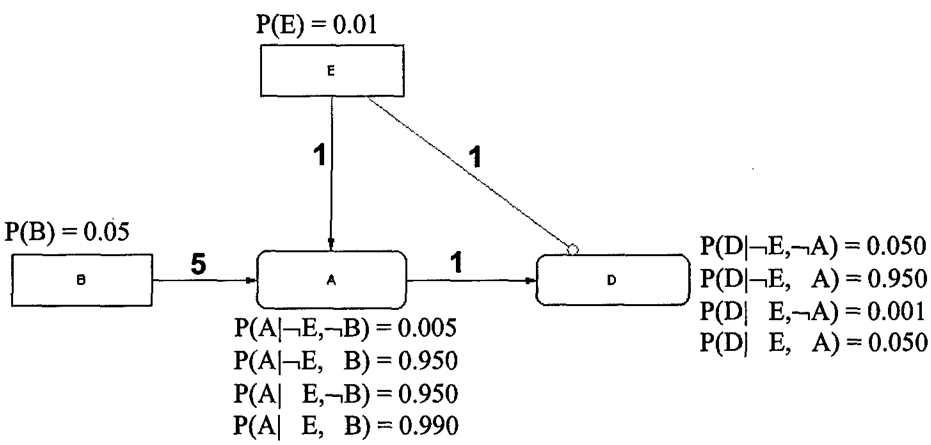

Figure 1: A Sample Influence Net

Formally, Influence Nets are Directed Acyclic Graphs (DAGs) where nodes in the graph represent random variables, while the edges between pairs of variables represent causal relationships. The following items characterize an IN:

1. A set of random variables that makes up the nodes of an IN. All the variables in the IN have binary states.
2. A set of directed links that connect pairs of nodes.
3. Each link has associated with it a pair of CAST Logic parameters that shows the causal strength of the link (usually denoted as g and h values).
4. Each non-root node has an associated CAST Logic parameter (denoted as the baseline probability), while a prior probability is associated with each root node.

Definition 1: An *Influence Net* is a tuple **(V, E, C, B)** where

V: set of Nodes,

E: set of Edges,

C represents causal strengths:

$\mathbf{E} \rightarrow \{ (\mathbf{h}, \mathbf{g}) \text{ such that } -1 < \mathbf{h}, \mathbf{g} < 1 \}$,

B represents a Baseline or Prior probability:

$\mathbf{V} \rightarrow [0,1]$

5.2.1 CAST Logic

Chang et al. [Chang et al., 1994] developed a formalism, at George Mason University, called CAusal STrength (CAST) logic as an intuitive and approximate language to elicit the large number of conditional probabilities from a small set of user-defined parameters. The logic has its roots in Noisy-OR approach [Agosta, 1991;Drudzel and Henrion, 1993]. In fact, it can be shown that the Noisy-OR approach is a special case of the CAST logic. The logic requires only a pair of parameter values for each dependency relationship between any two random variables. The logic is briefly explained with the help of an example shown in Figure 2. Readers interested in a detailed description of the CAST logic should refer to [Chang et al., 1994;Rosen and Smith, 1996].

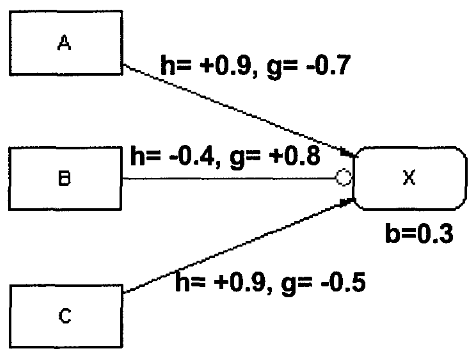

Figure 2: An Influence Network with CAST Logic Parameters

Figure 2 contains four nodes A, B, C and X. On each arc, two causal strengths are specified. These numbers represent the probability that a specified state of a parent node will cause a certain state in the child node. Positive values on arcs are causal influences that cause a node to occur with some probability, while negative values are influences that cause the negation of a node to occur with some probability. For instance, the arc between B and X has values –0.4 and 0.8. The first value, referred to as h, states that if B is true, then this will cause X to be false with probability 0.4, while the second value, referred to as g, states that if B is false, then this will cause X to be true with probability 0.8. Both h and g can take values in the interval (-1, 1). All non-root nodes are assigned a baseline probability, which is similar to the "leak" probability in the Noisy-OR approach. This probability is the user-assigned assessment that the event would occur independently of the modeled influences in a net.

There are four major steps in the CAST logic algorithm that converts the user-defined parameters into conditional probabilities:

a) Aggregate positive causal strengths
b) Aggregate negative causal strengths
c) Combine the positive and negative causal strengths, and
d) Derive conditional probabilities

In Figure 2, there are eight conditional probabilities that need to be computed to obtain the marginal probability of X. The marginal probability of X is computed as:

$$
\begin{aligned}
P(X) = {} & P(X \mid \neg A, \neg B, \neg C)\, P(\neg A, \neg B, \neg C) + P(X \mid \neg A, \neg B, C)\, P(\neg A, \neg B, C) \\
& + P(X \mid \neg A, B, \neg C)\, P(\neg A, B, \neg C) \quad + P(X \mid \neg A, B, C)\, P(\neg A, B, C) \\
& + P(X \mid A, \neg B, \neg C)\, P(A, \neg B, \neg C) \quad + P(X \mid A, \neg B, C)\, P(\neg A, \neg B, C) \\
& + P(X \mid A, B, \neg C)\, P(A, B, \neg C) \qquad + P(X \mid A, B, C)\, P(A, B, C)
\end{aligned}
$$
(1)

The four steps, described above, are used to calculate each of these eight conditional probabilities. For instance, to calculate the probability $P(X \mid A, B, \neg C)$, the h values on the arcs connecting A and B to X and the g value on the arc connecting C to X are considered. Hence, the set of causal strengths is {0.9, -0.4, -0.5}.

<u>Aggregate the Positive Causal Strengths:</u> In this step, the set of causal strengths with positive influence are combined. They are aggregated using the equation

$$PI = 1 - \prod_i (1 - C_i) \qquad \forall C_i > 0$$

where C_i is the corresponding g or h value having positive influence and PI is the combined positive causal strength. For our example

$$PI = 1 - (1 - 0.9) = 0.9$$

<u>Aggregate the Negative Causal Strengths:</u> In this step, the causal strengths with negative values are combined. The equation used for aggregation is

$$NI = 1 - \prod_i (1 - C_i) \qquad \forall C_i < 0$$

where C_i is the corresponding g or h value having negative influence and NI is the combined negative causal strength. Using the above equation, the aggregate negative influence is found to be:

$$NI = 1 - (1 - 0.4)\,(1 - 0.5) = 0.7$$

<u>Combine Positive and Negative Causal Strengths:</u> In this step, aggregated positive and negative influences are combined to obtain an overall net influence. The difference of these aggregated influences is taken. The overall influence is obtained by taking the ratio of this difference and the corresponding promoting or inhibiting influence. Mathematically,

$$\text{If } PI > NI$$
$$AI = \frac{PI - NI}{1 - NI}$$
$$\text{If } NI > PI$$
$$AI = \frac{NI - PI}{1 - PI}$$

Thus, the overall influence for the current example is

$$AI = (0.9 - 0.7) / (1 - 0.7) = .66$$

Derive Conditional Probabilities: In the final step, the overall influence is used to compute the conditional probability value of a child for the given combination of parents.

$$P(\text{child} \mid j\text{th state of parent states}) = \text{baseline} + (1 - \text{baseline}) \times AI \quad \text{when } PI > NI$$
$$= \text{baseline} - \text{baseline} \times AI \quad \text{when } PI < NI$$

Using the above equation, $P(X \mid A, B, \neg C)$ is obtained as:

$$P(X \mid A, B, \neg C) = 0.5 + 0.5 * 0.66 = .863$$

The steps explained above are repeated for the remaining seven conditional probabilities in Equation 1. If the experts had sufficient time and knowledge of the influences, then the probability matrix for each node can be used instead of g and h values. Also, after estimating the conditional probability matrix, if some entries do not satisfy the expert, then those entries can be modified and then used for computing the marginal probability of a node.

5.3 Timed Influence Nets

Influence Nets are designed to capture *static* interdependencies among variables in a system. However, a situation where the impact of a variable takes some time to reach the affected variable(s) cannot be modeled by either of the two approaches. Wagenhals et al. [1998] have added a special set of temporal constructs to the basic formalism of Influence Nets. The Influence Nets with these additional temporal constructs are called Timed Influence Nets (TINs) [Haider and Levis, 2004;Haider and Zaidi, 2004]. The temporal constructs allow a system modeler to specify delays associated with nodes and arcs. These delays may represent the information processing and communication delays present in a given situation. For example, in Figure 1, the inscription associated with each arc shows the corresponding time delay it takes for a parent node to influence a child node. For instance, event B influences the occurrence of event A in 5 time units.

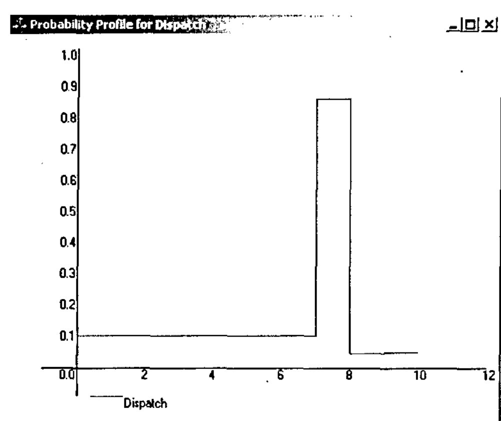

Figure 3: Probability Profiles of Event D

TINs have been experimentally used in the area of Effects Based Operations (EBOs) for evaluating alternate courses of actions and their effectiveness to mission objectives. [Wagenhals and Levis, 2000; Wagenhals and Levis, 2001; Wagenhals et al., 2003] The purpose of building a TIN is to evaluate and compare the performance of alternative courses of action. The impact of a selected course of action on the desired effect is analyzed with the help of a *probability profile*. Consider the net shown in Figure 1. Suppose it is decided that actions B and E are taken at time 1 and 7, respectively. Because of the propagation delay associated with each arc, the influences of these actions impact event D over a period of time. As a result, the probability of D changes at a different time instants. A probability profile draws these probabilities against the corresponding time line. The probability profile of event D is shown in Figure 3.

The following items characterize a TIN:

1. A set of random variables that makes up the nodes of a TIN. All the variables in the TIN have binary states.
2. A set of directed links that connect pairs of nodes.
3. Each link has associated with it a pair of parameters that shows the causal strength of the link (usually denoted as g and h values).
4. Each non-root node has an associated baseline probability, while a prior probability is associated with each root node.
5. Each link has a corresponding delay d (where $d \geq 0$) that represents the communication delay.
6. Each node has a corresponding delay e (where $e \geq 0$) that represents the information processing delay.
7. A pair (p, t) for each root node, where p is a list of real numbers representing probability values. For each probability value, a corresponding time interval is defined in t. In general, (p, t) is defined as

   $([p_1, p_2, \ldots, p_n], [[t_{11}, t_{12}], [t_{21}, t_{22}], \ldots, [t_{n1}, t_{n2}]])$,

   where $t_{i1} < t_{i2}$ and $t_{ij} > 0 \ \forall\ i = 1, 2, \ldots, n$ and $j = 1, 2$

The last item in the above list is referred to as an input scenario, or sometimes (informally) as a course of action. Formally, a TIN is described by the following definition.

Definition 2 Timed Influence Net (TIN)

A TIN is a tuple **(V, E, C, B, D_E, D_V, A)** where

V: set of Nodes,

E: set of Edges,

C represents causal strengths:

E $\rightarrow$ { (**h**, g) such that $-1 < \mathbf{h}, \mathbf{g} < 1$ },

B represents Baseline / Prior probability: **V** $\rightarrow$ [0,1],

$\mathbf{D_E}$ represents Delays on Edges: **E** $\rightarrow$ $\mathbf{Z}^+$, (where $\mathbf{Z}^+$ represent the set of positive integers)

$\mathbf{D_V}$ represents Delays on Nodes: **V** $\rightarrow$ $\mathbf{Z}^+$, and

A (input scenario) represents the probabilities associated with the state of actions and the time associated with them.

A: $\mathbf{R} \rightarrow \{([p_1, p_2, \ldots, p_n],[[t_{11},t_{12}], [t_{21},t_{22}], \ldots,[t_{n1},t_{n2}]])$

such that $p_i = [0, 1]$, $t_{ij} \rightarrow \mathbf{Z}^*$ and $t_{i1} \leq t_{i2}$, $\forall\ i = 1, 2, \ldots, n$ and $j = 1, 2$

where $\mathbf{R} \subset \mathbf{V}$ }(where $\mathbf{Z}^*$ represent the set of nonzero positive integers)

5.4. Limitations of TIN and the Proposed Enhancements

5.4.1 Modeling of Memory

The existing TINs are not capable of modeling the impact of different sequences of actions on the desired effect. This behavior is because of the underlying assumption in TINs that events are *memoryless*, i.e. the probability of occurrence of an event at a particular time instant does not depend upon its own probabilities of occurrence during the previous time instants. As a consequence, the probability of an event depends only upon the actions executed so far and not on the sequence in which these actions are executed. The proposed approach adds an optional *self-loop* to each node. The events having self-loops are no longer assumed to be memoryless. Like other arcs in a TIN, a self-loop is also specified using the CAST logic. A higher value (either positive or negative) of the parameters imitates strong memory while a lower value imitates weak memory. If both parameters (h and g) are set to zero then this is equivalent of having no self-loop. Thus, this class of TINs is a superset of the TINs that were defined in Definition 2. In the TIN of Figure 1, if events A and D depend upon their previous states, then this phenomenon is captured by adding a self-loop to each of them as shown in Figure 4.

The addition of self-loop not only changes the final probability of the variable of interest, but it also has an affect on the trajectory of the probability profile. Consider the TIN shown in Figure 5. It has three variables A, B, and C. In the absence of a self-loop, the probability of event C depends only upon the probability of its parents, that is, A and B. Suppose two courses of action required to be evaluated for this model. In the first course of action (COA 1), actions A and B are taken at times 10 and 12, respectively while in the second course of action 2 (COA 2), A and B are taken at time 12 and 10, respectively. The respective probability profiles of C as a result of these courses of action are shown in Figure 6(a) and 6(b). Despite the fact that the trajectories shown in the two profiles differ significantly, the final probability of event C is same (0.85) in both profiles. This behavior is due to the fact that the underlying TIN model is memoryless. Thus, no matter what the sequence of actions A and B is, the likelihood of occurrence of C is same once both actions are taken.

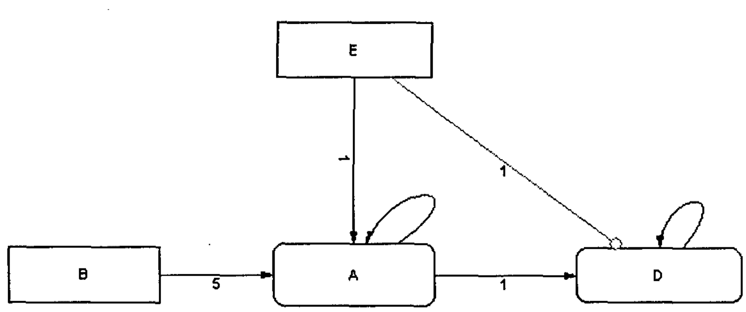

Figure 4: A Timed Influence Net with Self-loop

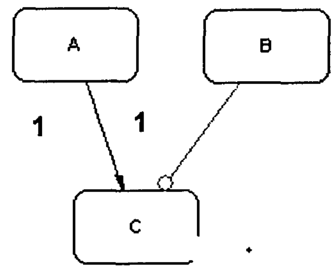

Figure 5: A TIN having 3 Nodes

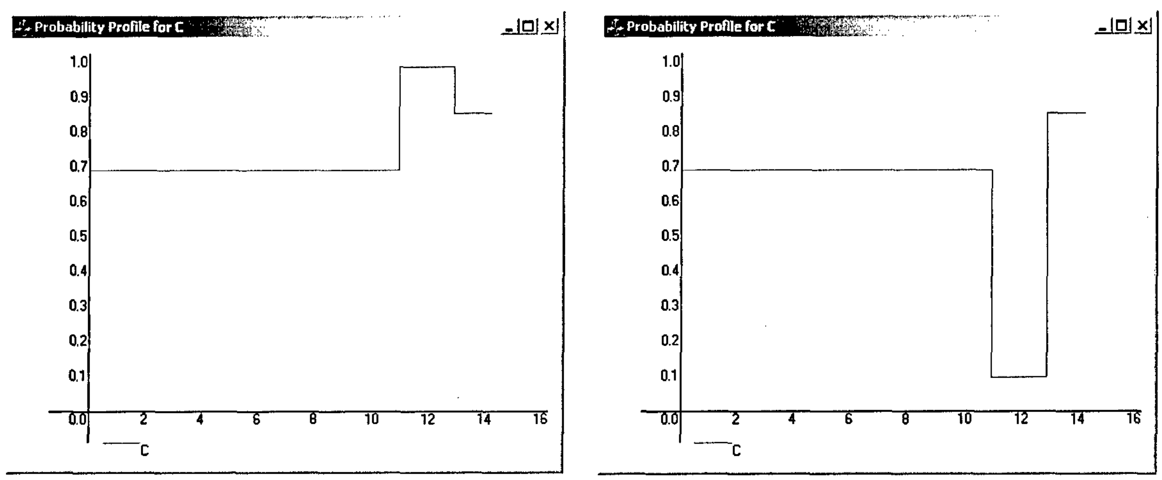

(a) COA 1

(b) COA 2

Figure 6: Probability Profiles of Event C in the TIN of Figure 5

In contrast to the given situation, suppose the likelihood of C at a particular time instance depends upon its own likelihood in the past. The proposed methodology attempts to model this situation by adding a self-loop to event C. The modified TIN is shown in Figure 7. The text associated with the self-loop shows the corresponding CAST logic parameters. In addition to their normal semantics, the parameters attached to a self-loop also represent the strength of the memory associated with the corresponding variable. For instance, high values of g and h strongly causes a node to remain in its previous state, while a lower values of g and h represents a weak memory and thus the previous state of a variable does not have a large influence on its current state. The two courses of action described earlier (COA 1 and COA 2) are executed for the model of Figure 7(a) and the respective profiles are shown in Figure 8.

It can be seen from the profiles that the final probability of event C is different in the two profiles. This change in the behavior of the TIN occurs because of the fact that now the present likelihood of C depends upon its likelihood in the past along with the probabilities of its parents. For instance, in the profile of Figure 8(a), event A happens first which causes an increase in the probability of C (0.85) as the occurrence of A has a strong positive influence on the occurrence of C. B happens after A. Despite its negative influence, B fails to decrease the likelihood of C as C has a strong memory that causes it to remain in the previous state along with the fact that a strong positive influence from A counterbalances a moderate negative influence from B. Thus, the final probability of C is 0.87. In the second profile (Figure 8(b)), B happens first and due to its negative influence on C the probability of C is decreased to 0.56. A happens next and it slightly increases the probability of C to 0.63 but not as much as it is increased in COA 1 because of the dependency of C on its previous state. While computing the profiles of Figures 6 and 8, it can be noticed that the profiles have quite a different behavior in both courses of action.

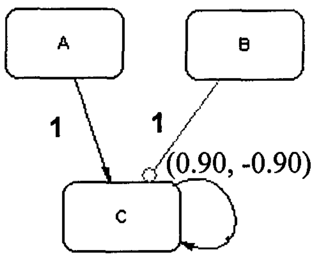

(a) Strong Memory

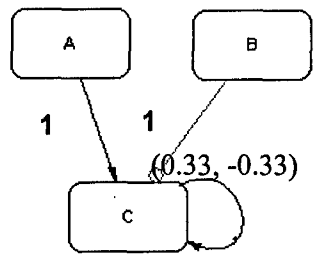

(b) Weak Memory

Figure 7: Different Levels of Memory Modeled Using Self-Loops

As mentioned earlier, if the h and g values associated with a self-loop are low, then the loop represents a weak influence of the previous state of a node on its current state. Suppose in the model of Figure 7(a), the g and h values associated with the self-loop are revised and are as shown in Figure 7(b). The same two courses of action (COA 1 and COA 2) are executed in this situation and the resultant probability profiles are shown in Figure 9: a weak memory has resulted in the final probabilities very close to what is obtained in the profiles based on a memoryless TIN (Figure 6).

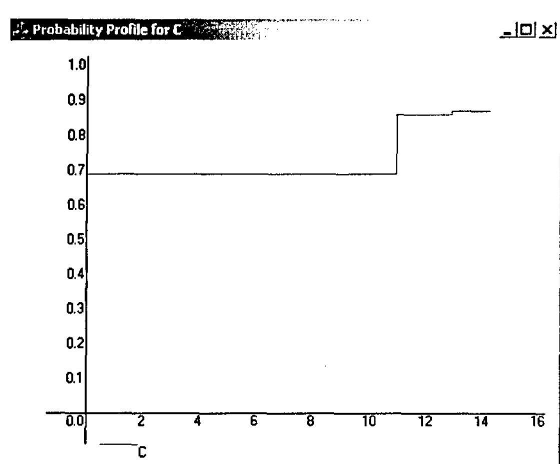

(a) COA 1

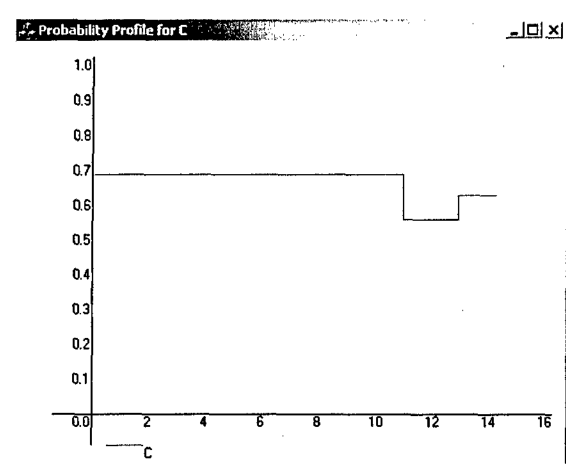

(b) COA 2

Figure 8: Probability Profiles of Event C in the TIN of Figure 7(a)

Up until now, it is assumed that a node's likelihood at a previous time stamp is used to update its current likelihood when a new piece of information arrives from one of its parents. A self-loop can also be used to update the probability of a node at a regular time interval. This time interval is specified as the delay associated with a self-loop. Thus a self-loop can be used to model *decay* in the belief of a node as the time passes and no new information from its parents influences it. Suppose in the model of Figure 7, the self-loop associated with node C has a delay of 1 time unit which means that the probability of C is updated after every 1 time unit regardless of whether there is new information coming from its parents or not. In the sequel, if the delay associated with a self-loop has a value of zero then it means that a previous value of a node is used to update its current likelihood only when there is new information coming from its parents. Positive values other than zero indicate that the update would occur at a regular time interval.

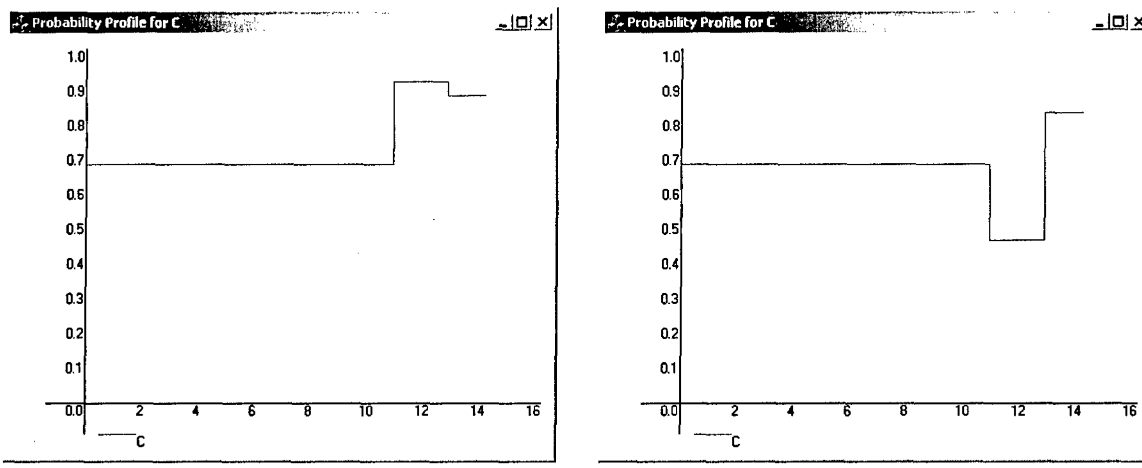

(a) COA 1 (b) COA 2

Figure 9: Probability Profiles of Event C in the TIN of Figure 7(b)

5.4.2 Modeling of Time Varying Influences

Events happen and they influence other relevant events. In many cases, the intensity of their influences decays over time. Thus, an event having a very strong influence at the time of its occurrence on another event might have an insignificant influence after a certain period of time. In other words, the influence of an event is *time-variant*. The *time-varying* property also holds true for the state of an action. An action may occur in two different states during two different time intervals. In TINs terminology, these two types of time-varying properties are referred to as *persistence*. The one related to the time dependent influence of an action is called *persistence of influence,* while the one related to the time-dependent state of an action is called *persistence of action*. Among these two types of persistence, a TIN currently models the latter one only. It assumes that the causal strength of the influences does not change over time, i.e., the underlying stochastic model is *stationary*. Thus, it lacks the ability to model persistence of influence. This paper attempts to overcome this limitation of TINs. The proposed approach enables a system modeler to model *non-stationary* influences. Instead of asking a modeler to specify single- valued influences, the proposed approach would allow the modeler to specify various strengths of influences and their corresponding window of effectiveness.

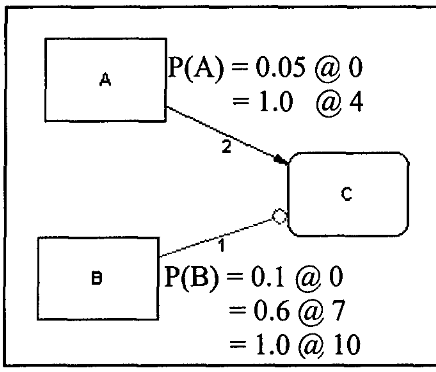

Influence of A on C when information at A is t time units old
Strong: $2 \leq t < 4$
Moderate: $4 \leq t \leq 6$
Low: $t > 6$
Influence of B on C when information at B is t time units old
Strong: $1 \leq t < 3$
Low: $t > 3$

Figure 10: A TIN Having Time-Variant Influences

Consider the TIN of Figure 10. The prior probability of nodes A and B at time 0 is 0.05 and 0.1, respectively. Action A is taken at time 4 while the probability of occurrence of B becomes 0.6 at time 7 and 1 at time 10. The CAST logic parameters associated with the arc are time-varying and are read in the following manner. A has a high positive influence on C, if the change occurred at A is 2 to 3 time units old, and its influence is moderate, if the change occurred at A is 4 to 5 time units, while its influence is low if the change occurred at A is more than 6 time units old. For simplicity "high influence" is assumed to mean that both h and g have the same values though with opposite signs (one is positive and the other is negative). Similarly, B has a strong negative influence on C when the change that occurred at B is 1 to 2 time units old, while it has a low influence when the change occurred at B is more than 2 time units old. Due to the provided input scenario, the probability of C is updated at time 6, 8, and 11 as the time delays between A and C and B and C are 2 and 1, respectively. C is updated at time 6 because action A is taken at time stamp 4. The last change that occurred at B is at time 0. Thus, the probability of B used in computing the marginal probability of C is 0.1. Since this value is 6 time units old, while computing the CPT values for node C a low negative influence of B on C is considered. The TIN with a particular instance of the CAST logic parameters along with the prior probabilities is shown in Figure 11. The Conditional Probability Table (CPT) values computed under this situation are also shown beside node C. Based on the parameters shown in the figure, the marginal probability of C at time 6 is computed as given below.

$$
\begin{aligned}
P(C) = {} & P(C|\neg A,\neg B) \times P(\neg A) \times P(\neg B) + P(C|\neg A,\ B) \times P(A) \times P(B) \\
& + P(C|\neg A,\ B) \times P(\neg A) \times P(\ B) + P(C|\ A,\ B) \times P(A) \times P(B) \\
= {} & 0.07 \times 0 \times 0.9 + 0.03 \times 0 \times 0.9 + 0.97 \times 1 \times 0.1 + 0.93 \times 1 \times 0.9 \\
= {} & 0.93
\end{aligned}
$$

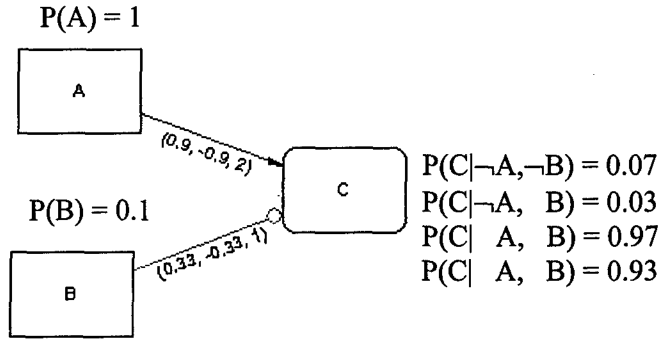

Figure 11: An Instance of the TIN of Figure 10

The next update of P(C) occurs at time 8. At this time instance, the marginal probability of A is 4 time units old; thus a moderate positive influence of A on C is considered while computing the CPT values. The probability of B is only 2 time units old and has a strong negative influence on C. The resultant parameters, along with the CPT values, are shown in Figure 12. The probability of C at time 8 is computed as shown below.

$$P(C) = P(C|\neg A,\neg B) \times P(\neg A) \times P(\neg B) + P(C|\neg A,\ B) \times P(A) \times P(B)$$
$$+ P(C|\neg A,\ B) \times P(\neg A) \times P(\ B) + P(C|\ A,\ B) \times P(A) \times P(B)$$

$$= 0.85 \times 0 \times 0.4 + 0.02 \times 0 \times 0.6 + 0.98 \times 1 \times 0.4 + 0.15 \times 1 \times 0.6$$
$$= 0.48$$

The last update of P(C) occurs at time 11. The marginal probability of A is 7 time units old, while B's is 2 time units old. Thus, a low positive influence from A and a strong negative influence from B are considered. The updated probability of C is found to be 0.07. The above analysis demonstrated how non-stationary CAST logic parameters have resulted in non-stationary CPT values that are used in computing the probability of C at various time stamps. Thus, despite the fact that an action is still in effect, it may loose its significance as time passes by. The non-stationary CPT values used in the above computations are compared in Table 1 along with the time of their computation.

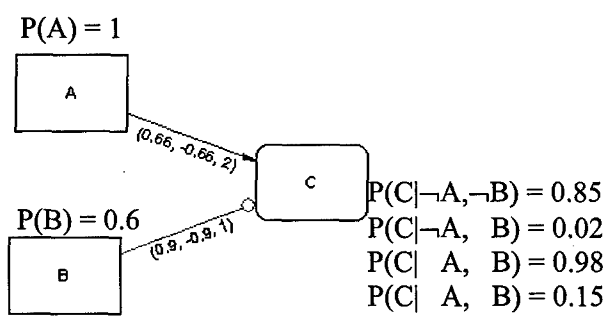

Figure 12: Another Instance of the TIN of Figure 10

Table1: for Non-Stationary CPTs

		Time	
Parents Combination	6	8	11
P(C|¬A,¬B)	0.07	0.85	0.93
P(C|¬A, B)	0.03	0.02	0.03
P(C| A,¬B)	0.97	0.98	0.97
P(C| A, B)	0.93	0.15	0.07

5.5. Dynamic Influence Nets (DINs)

The incorporation of the proposed structural and parametric changes in TINs, as described in the previous sections, would enable a system modeler to observe the impact of repeated actions. For instance, an air-strike on a bridge makes it inoperable for several days. The current implementation of TINs would assume that the influence of the air-strike remains the same throughout the campaign. It is obvious that the assumption is unrealistic. Furthermore, in the event of a new air-strike, a TIN would discard the impact of the previous air-strike as the events in a TIN are assumed to be memory-less. The proposed approach, which allows time-varying influences and incorporation of memory through self-loops, models this situation in a more intuitive manner. Like other arcs in a TIN, a self-loop also represents influence – from the previous state of a node to its current state. Thus, time-varying parameters can be associated with a self-loop too. For the air-strike example, the presence of self-loop would combine the influences of both (or many) air-strikes while the strength of the self-loop accounts for the time delay between the two air-strikes. If the timing of two air-strikes is far apart, then there is almost no influence of the first strike on the operability of the bridge (provided that the bridge has been rebuilt), but if the two strikes occurred very close in time then their impact would be more destructive. In other words the impact of two actions on the effect convolves. The issue is further explained with the help of the following example. Suppose the variables in the model of Figure 10 have the following descriptions:

A - Regional Countries Opposes Sanctions against Country R
B – Country G Threatens to Take Unilateral Actions against Country R
C – Leader of Country R Decides to Accept UN Demands

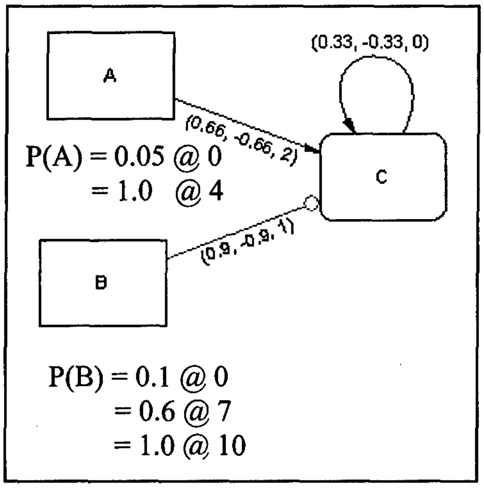

Influence of A on C when information at A is t time units old
Strong: $2 \leq t < 4$
Moderate: $4 \leq t \leq 6$
Low: $t > 6$

Influence of B on C when information at B is t time units old
Strong: $1 \leq t < 3$
Low: $t > 3$

Figure 13: A TIN with Self-Loop and Time-Varying Influences

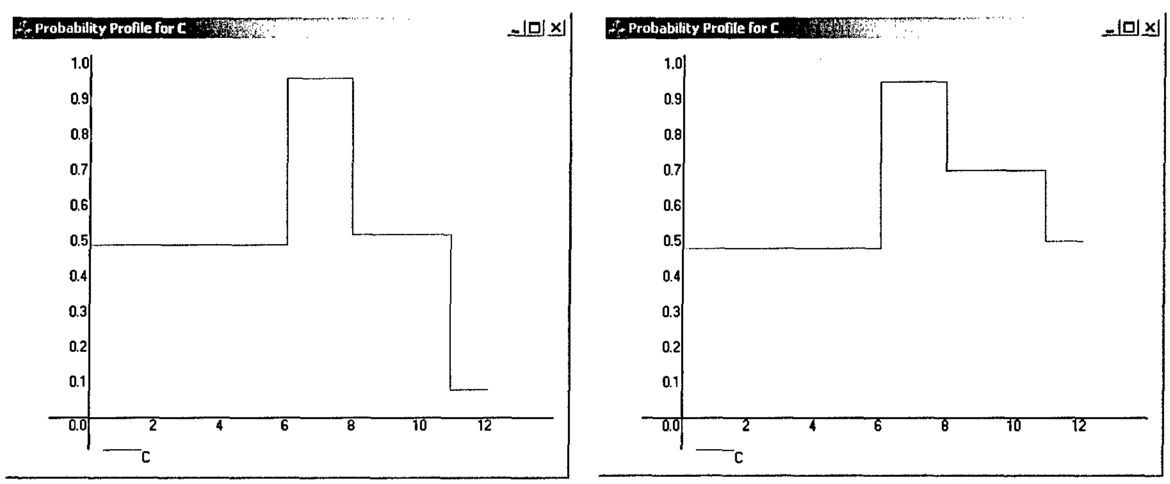

(a) Profile of event C in a DIN

(b) Profile of event C in a TIN

Figure 14: Comparison of Profiles Generated by a TIN and a DIN

Further assume that the belief of event C at a particular time depends upon its own belief at a previous time instance though not very strongly. This behavior is modeled by adding a self-loop having a low influence on event C. The resultant model is shown in Figure 13. The probabilities of actionable events A and B are changed at various time stamps, as described earlier and shown in the figure. The resultant probability profile of C is shown in Figure 14(a). If the same situation is modeled using an existing TIN, that is, without the self-loop and time-invariants CAST logic parameters, then the resultant profile of C is as shown in Figure 14(b). Currently, there is no validation technique that helps us in identifying which profile is a better representation of the situation at hand, but it can be said that the profile shown in Figure 14(a) is more in agreement with intuition than the profile of Figure 14(b). The impact of B is more dominating in the profile of Figure 14(a) as event A happened 6 time units earlier and has lost its significance. Furthermore, the previous state of event C also has an impact on its current state. Thus, a different sequence of actions would have resulted in a completely different outcome. Profile of Figure 14(b) fails to capture these characteristics.

The incorporation of the new constructs (self-loop and time-varying parameters) in the framework of TIN based modeling and reasoning enhances the capabilities of this modeling paradigm in terms of capturing dynamic uncertain scenarios. A TIN with these additional constructs has been defined as a Dynamic Influence Net (DIN). The following items characterize a DIN while a formal definition is given in Definition 3.

1. The nodes of a DIN are set of random variables. All the variables in the DIN have binary states.
2. A set of directed links that connect pairs of nodes. A node can also have an optional self-loop.
3. A pair (c, t) for each link, where c is a list of tuples representing the CAST logic parameters. For each element in c, a corresponding time interval is defined in t. This interval represents the time during which the corresponding element in c is in effect. In general, (c, t) is defined as
   $([(h_1, g_1) (h_2, g_2), \ldots, (h_n, g_n)], [(t_{11}, t_{12}), (t_{21}, t_{22}), \ldots, (t_{n1}, t_{n2})]$
   where $t_{i1} < t_{i2}$ and $t_{ij} > 0 \ \forall\ i = 1, 2, \ldots, n$ and $j = 1, 2$
4. Each non-root node has an associated baseline probability, while a prior probability is associated with each root node.
5. Each link has a corresponding delay d (where $d \geq 0$) that represents the communication delay.
6. Each node has a corresponding delay e (where $e \geq 0$) that represents the information processing delay.
7. A pair (p, t) for each root node, where p is a list of real numbers representing probability values. For each probability value, a corresponding time interval is defined in t. In general, (p, t) is defined as
   $([p_1, p_2, \ldots, p_n], [[t_{11}, t_{12}], [t_{21}, t_{22}], \ldots, [t_{n1}, t_{n2}]]$
   where $t_{i1} < t_{i2}$ and $t_{ij} > 0 \ \forall\ i = 1, 2, \ldots, n$ and $j = 1, 2$

Definition 3: A *Dynamic Influence Net* is a tuple **(V, E, C, B, $\mathbf{D_E}$, $\mathbf{D_V}$, A)** where

V: set of Nodes,

E: set of Edges,

C represents causal strengths:

$\mathbf{E} \rightarrow \{([(h_1, g_1) (h_2, g_2), \ldots, (h_n, g_n)], [(t_{11}, t_{12}), (t_{21}, t_{22}), \ldots, (t_{n1}, t_{n2})])$
such that $-1 < h_i, g_i < 1$ }, $t_{ij} \rightarrow \mathbf{Z}^+$ and $t_{i1} \leq t_{i2}, \forall\ i = 1, 2, \ldots, n$ and $j = 1, 2\}$

B represents Baseline / Prior probability: $\mathbf{V} \rightarrow [0,1]$,

$\mathbf{D_E}$ represents Delays on Edges: $\mathbf{E} \rightarrow \mathbf{Z}^+$,

$\mathbf{D_V}$ represents Delays on Nodes: $\mathbf{V} \rightarrow \mathbf{Z}^+$, and

A (input scenario) represents the probabilities associated with the set of actions and the time associated with them.

$\mathbf{A}: \mathbf{R} \rightarrow \{([p_1, p_2, \ldots, p_n], [[t_{11}, t_{12}], [t_{21}, t_{22}], \ldots, [t_{n1}, t_{n2}]])$
such that $p_i = [0, 1]$, $t_{ij} \rightarrow \mathbf{Z}^*$ and $t_{i1} \leq t_{i2}, \forall\ i = 1, 2, \ldots, n$ and $j = 1, 2$ where $R \subset V$ }

5.6 Conclusions

The paper presents structural and parametric enhancements to Timed Influence Nets based modeling framework. Currently, nodes in a TIN are considered memoryless. This inability results in lack of modeling the impact of different sequences of actions on a desired effect. TINs also fail to capture time-varying influences. The proposed enhancements would allow a system modeler to specify such influences. Thus, the modeler would be able to specify both stationary and non-stationary influences. Furthermore, the dependency of an event on its previous state could also be modeled by adding a self-loop to the corresponding node. The incorporation of self-loop adds memory to the existing memory-less TIN. The addition of both self-loop and time-varying influences would enable a modeler to model impacts of repeated actions on an effect. Currently, in the event of repeated actions, a TIN only considers the latest impact on the effect while ignoring the previous attempts. The proposed DIN would convolve the impact of repeated actions on the desired effect and, thus, further enhance the capabilities of TINs based modeling paradigm.

References

Agosta, J. M., "Conditional Inter-Causally Independent Node Distributions, a Property of Noisy-or Models," Proceedings of 7th Conference on Uncertainty in Artificial Intelligence, 1991.

Chang, K. C., Lehner, P. E., Levis, A. H., Zaidi, S. A. K. and Zhao, X., "On Causal Influence Logic," George Mason University, Technical Report Dec. 1994.

Drudzel, M. J. and Henrion, M., "Intercausal Reasoning with Uninstantiated Ancestor Nodes," Proceedings of the 9th Conference on Uncertainty in AI, 1993.

Haider, S. and Levis, A. H., "A Heuristic Approach for Belief Updating in Timed Influence Nets," Proceedings of 2004 Command and Control Research and Technology Symposium, San Deigo, 2004.

Haider, S. and Zaidi, A. K., "Transforming Timed Influence Nets into Time Sliced Bayesian Networks," Proceedings of 2004 Command and Control Research and Technology Symposium, San Diego, 2004.

Kschischang, F. R. and Frey, B. J., "Iterative Decoding of Compound Codes by Probability Propagation in Graphical Models," IEEE Journal on Selected Areas in Communication, vol. 16, pp. 219-230, 1998.

McEliece, R. J., MacKay, D. J. and Cheng, J. F., "Turbo Decoding as an Instance of Pearl's Belief Propagation Algorithm," IEEE Journal on Selected Areas in Communication, vol. 16, pp. 140--151, 1998.

Murphy, K. P., Weiss, Y. and Jordan, M. I., "Loopy Belief Propagation for Approximate Inference: An Empirical Study," Proceedings of 15th Conference on Uncertainty in Artificial Intelligence, 1999.

Pearl, J., Probabilistic Reasoning in Intelligent Systems: Network of Plausible Inference: Morgan Kaufmann, 1987.

Rosen, J. A. and Smith, W. L., "Influence Net Modeling with Causal Strengths: An Evolutionary Approach," Proceedings of the 1996 Command and Control Research and Technology Symposium, Monterey CA, 1996.

Wagenhals, L. W. and Levis, A. H., " Modeling Support of Effect-Based Operations in War Games," Proceedings of the 2002 Command and Control Research and Technology Symposium, 2002.

Wagenhals, L. W., Levis, A. H. and McCrabb, M. B., "Effects-Based Operations; a Historical Perspective of a Way Ahead," Proceedings of 8th International Command and Control Research and Technology Symposium, 2003.

Wagenhals, L. W. and Wentz, L. K., "New Effects-Based Operations Models in War Games," Proceedings of 8th International Command and Control Research and Technology Symposium, Washington DC, 2003.

SECTION 6

On Finding Effective Courses of Action in a Complex Situation Using Evolutionary Algorithms

Sajjad Haider and Alexander H. Levis

6.1. Introduction

In this modern world, any organization's, either military or business, effort is guided by a set of objectives/desired effects that it plans to achieve in an uncertain complex environment. A decision maker, working in that organization, is typically confronted by the task of finding a supposedly optimal strategy to achieve these effects. The pre-requisite of this task is the modeling of cause-effect/relevance relationships among the variables that exist in the environment. The last two decades have seen an emergence of the use of probabilistic reasoning framework as a modeling tool for capturing such relationships. Commonly referred to as Bayesian Networks (BN) [Charniak, 1991; Neapolitan, 2003; Pearl, 1987], the framework uses a graph-theoretic representation to explicitly show the dependencies among variables in a problem domain. Despite their ability to represent uncertain domain in a compact and easy to read manner, the knowledge acquisition process and inference mechanism of BNs is intractable. [Cooper, 1990; Dagum and Luby, 1993] Influence Nets (INs), an instance of the Bayesian framework, were proposed [Chang et al., 1994; Rosen and Smith, 1996] a decade ago to overcome the intractability issues present in BNs. They employ an approximation inference algorithm and non-probabilistic knowledge acquisition interface that does not require an exponential number of parameters.

Both Bayesian Networks and Influence Nets are designed to capture static interdependencies among variables in a system. However, a situation where the impact of a variable takes some time to reach the affected variable(s) cannot be modeled by either of the two approaches. In the last several years, efforts have been made to integrate the notion of time and uncertainty. [Boyen and Koller, 1998; Hanks et al., 1995; Kjaerulff, 1992; Santos and Young, 1999] Wagenhals et al. [Wagenhals et al. 1998] have added a special set of temporal constructs to the basic formalism of Influence Nets. The Influence Nets with these additional temporal constructs are called Timed Influence Nets (TINs) [Haider and Levis, 2004; Haider and Zaidi, 2004]. TINs have been experimentally used in the area of Effects Based Operations (EBOs) for evaluating alternate courses of actions and their effectiveness to mission objectives. [Wagenhals and Levis, 2000; Wagenhals and Levis, 2001; Wagenhals et al., 2003].

Once a dynamic uncertain situation is modeled using TIN; a decision maker is interested in the identification of a set(s) of actions and their time of execution that would maximize the likelihood of achieving the desired effect. The task is sometimes informally referred to as the best course of action (COA) determination. This paper applies Evolutionary Algorithms (EA) to accomplish the task of effective COA determination. It is important to note that any approach, which attempts to automate the process of identification of an effective COA, has to consider several temporal and causal constraints that exist among actionable events. Currently, TINs do not have a mechanism to specify these constraints. The paper proposes a constraint specification language (CSL) that aids a system modeler in specifying these temporal and causal constraints. The proposed technique takes these constraints into consideration while identifying an effective strategy in a given situation.

The optimization problem described above belongs to the category of Mixed Integer Nonlinearly Constraint Optimization. The nonlinearities are encoded in the form of conditional probabilities, while integral constraints result from the binary nature of actionable events. Besides providing the optimal or near optimal courses of actions, the proposed approach also provides a scheme to generalize the alternate effective courses of actions. It should be mentioned that, from an academic point of view, the problems presented in this paper are relatively new and not much work has been reported in the literature that attempts to solve them in an automated fashion. A few exceptions are the work presented by Wagenhals [Wagenhals, 2000] and Haiying et al. [Haiying et al., 2004]. The current practice among the community is to use hit and trial methods to identify an effective course of action.

The rest of the paper is organized as follows. Sections 6.2 and 6.3 provide a brief overview of Timed Influence Nets and Evolutionary Algorithms, respectively. Section 6.4 describes the issues that need to be considered while finding an effective course of action. The section also explains the proposed constraint specification language and specifics of the EA used in this paper. The results produced by the proposed approach, when applied on a real model, are explained in Section 6.5. Finally Section 6.6 concludes the paper and provides future research directions.

6.2 Timed Influence Nets

The modeling of the causal relationships in TINs is accomplished by creating a series of cause and effect relationships between some desired effects and the set of actions that might impact their occurrence in the form of an acyclic graph. The actionable events in a TIN are drawn as root nodes (nodes without incoming edges). A desired effect, or an objective in which a decision maker is interested, is modeled as a leaf node (node without outgoing edges). Typically, the root nodes are drawn as rectangles while the non-root nodes are drawn as rounded rectangles. Figure 1 shows a partially specified TIN. Nodes B and E represent the actionable events (root nodes) while node C represents the objective node (leaf node). The directed edge with an arrowhead between two nodes shows the parent node promoting the chances of a child node being true, while the roundhead edge shows the parent node inhibiting the chances of a child node being true. The inscription associated with each arc shows the corresponding time delay it takes for a parent node to influence a child node. For instance, event B, in Figure 1, influences the occurrence of event A after 5 time units.

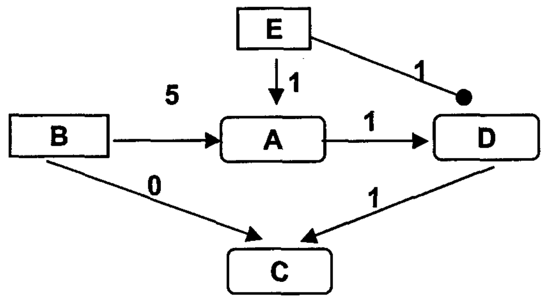

Figure 1: An Example Timed Influence Net (TIN)

The purpose of building a TIN is to evaluate and compare the performance of alternative courses of actions. The impact of a selected course of action on the desired effect is analyzed with the help of a *probability profile*. Consider the TIN shown in Figure 1. Suppose the following *input scenario* is decided: actions B and E are taken at times 1 and 7, respectively. Because of the propagation delay associated with each arc, the influences of these actions impact event C over a period of time. As a result, the probability of C changes at different time instants. A probability profile draws these probabilities against the corresponding time line. The probability profile of event C is shown in Figure 2.

The following items characterize a TIN:

1. A set of random variables that makes up the nodes of a TIN. All the variables in the TIN have binary states.
2. A set of directed links that connect pairs of nodes.
3. Each link has associated with it a pair of parameters that shows the causal strength of the link (usually denoted as g and h values).
4. Each non-root node has an associated baseline probability, while a prior probability is associated with each root node.
5. Each link has a corresponding delay d (where $d \geq 0$) that represents the communication delay.
6. Each node has a corresponding delay e (where $e \geq 0$) that represents the information processing delay.
7. A pair (p, t) for each root node, where p is a list of real numbers representing probability values. For each probability value, a corresponding time interval is defined in t. In general, (p, t) is defined as

   $([p_1, p_2, \ldots, p_n], [[t_{11}, t_{12}], [t_{21}, t_{22}], \ldots, [t_{n1}, t_{n2}]])$,

   where $t_{i1} < t_{i2}$ and $t_{ij} > 0 \ \forall\ i = 1, 2, \ldots, n$ and $j = 1, 2$

The last item in the above list is referred to as input scenario, or sometimes (informally) as course of action. Formally, a TIN is described by the following definition.

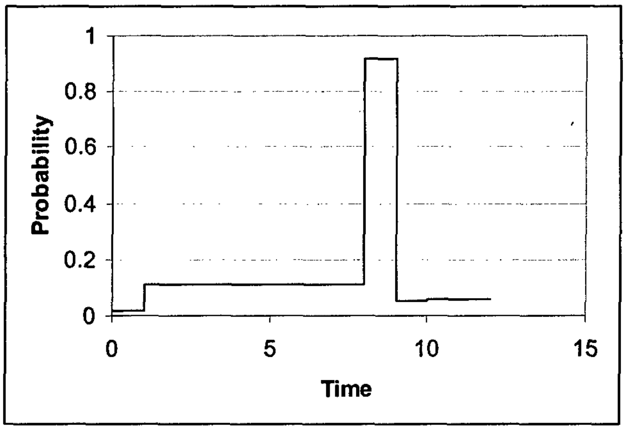

Figure 2: Probability Profile for Node C

Definition 1 Timed Influence Net (TIN)

A TIN is a tuple (**V, E, C, B, D_E, D_V, A**) where

V: set of Nodes,

E: set of Edges,

C represents causal strengths:

$\mathbf{E} \rightarrow \{ (\mathbf{h}, \mathbf{g}) \text{ such that } -1 < \mathbf{h}, \mathbf{g} < 1 \}$,

B represents Baseline / Prior probability: $\mathbf{V} \rightarrow [0,1]$,

$\mathbf{D_E}$ represents Delays on Edges: $\mathbf{E} \rightarrow \mathbf{N}$,

$\mathbf{D_V}$ represents Delays on Nodes: $\mathbf{V} \rightarrow \mathbf{N}$, and

A (input scenario) represents the probabilities associated with the state of actions and the time associated with them.

A: $\mathbf{R} \rightarrow \{([p_1, p_2, \ldots, p_n],[[t_{11},t_{12}], [t_{21},t_{22}], \ldots,[t_{n1},t_{n2}]])$

such that $p_i = [0, 1]$, $t_{ij} \rightarrow \mathbf{Z}$ and $t_{i1} \leq t_{i2}$, $\forall$ i = 1, 2, …., n and j = 1, 2 where

$\mathbf{R} \subset \mathbf{V} \}$

6.3 Evolutionary Algorithm

Evolutionary Algorithms (EAs) can be interpreted as a parallel adaptive search procedure. They have been applied to a wide variety of application areas including optimization, search, learning, automated programming, adaptation, etc. They are modeled after the organic evolutionary processes found in nature. The current EAs have their roots in three distinct efforts that were initiated in parallel almost four decade ago: Evolution Strategies [Schwefel, 1975;Schwefel, 1995] Evolutionary Programming [Fogel et al., 1966], and Genetic Algorithms [DeJong, 1975;Goldberg, 1989].

An EA consists of a *population* of *individual solutions* that are selected and modified in order to discover overall better solutions in the search space. An individual in the population is referred to as *genome* or *phenome*. One (or more) solution(s) is (are) modified to produce new *offspring* in the population. Each solution is evaluated using a *fitness* value that represents the quality of the solution in the context of a given problem. The major steps involved in a canonical EA are shown in Table 1 [DeJong, 2005]. There are three major design issues, namely, representation, selection, and variation operators, need to be considered while designing an EA for a specific problem. These design issues are briefly explained in the following sub-sections.

Table 1: A Canonical Evolutionary Algorithm

Randomly generate the initial population
Do until a stopping criterion is met
    Select parent(s) using a *selection procedure*.
    Create new offspring(s) by applying the *variation operators* on the parents.
    Compute the fitness of the new offspring(s).
    Select member(s) of the population to die using a selection procedure.

6.3.1 Representation

Representation is among the most critical design issues while developing an EA. In an EA, there are two possible types of representations. When a representation is stated in its natural state it is said to be a *phenotype*. For example, a real valued optimization problem would be represented as a set of real valued coordinates. On the other hand, it may be necessary to map a phenotype representation into another structure to make it easier for the algorithm to modify and exchange information. In many cases the phenotype can be mapped into a representation that resembles, at a very high level, a sequence of local structures or building blocks. In this case, the representation is referred to as a *genotype*.

6.3.2 Selection

Selection in evolutionary algorithms is the process of choosing which individuals reproduce offspring and which individuals survive to the next generation. When selection is used to choose which individuals reproduce, the process is referred to as pre-selection (parent(s) selection). When it is used to select the individuals that survive to the next generation it is called post-selection (survival selection). Selection can further be categorized as deterministic or probabilistic. Deterministic selection tends to behave more like greedy hill-climbing algorithms and exploit the nearest areas with promising solutions. Probabilistic selection schemes are more exploratory and search the landscape.

An important decision required when deciding on what type of selection schemes to use is whether to emphasize *exploration* or *exploitation*. Schemes based on exploration, are said to have a low selection pressure, while schemes based on exploitation, are said to have greater selection pressure. It can be said that the selection pressure is a vague measure of how often more fit individuals are selected to reproduce and/or live to the next generation.

It should be mentioned that EAs are not necessarily good at finding optimal solutions. For complex problems, the global optimal solution may be very difficult to locate for any algorithm. This leads to a trade-off between exploring the landscape for areas that appear to hold good solutions and exploiting the good area found. This can be a very difficult balancing act. Too much exploration may result in not finding the local optima in the regions explored, while too much exploitation can lead to behavior similar to a greedy hill climbing algorithm. The key is to find a selection pressure that balances exploration and exploitation.

Selection schemes can be further categorized into generational or steady-state schemes. A selection scheme is generational when the entire current population is replaced by its offspring to create the next generation; while, a scheme is referred to as steady-state when a selected few offspring replace a few members of the current generation to form the next generation. Some of the popular selection schemes are given below:

<u>Fitness Proportional:</u> In this scheme, individuals are selected based on their fitness in proportion to the other individuals in the population.

<u>Rank Selection:</u> In rank selection, individuals in the population are first ranked by fitness and then selected for reproduction based on a probability proportional to rank.

<u>Binary Tournament:</u> In this scheme, two individuals are randomly selected from the population and compared. The one with the highest fitness is selected for reproduction. Then another two individuals are randomly selected and the best fit kept as the mate to the first parent.

<u>Truncation:</u> This scheme is used during post-selection. Populations of parents and offspring are merged together and top k fittest individuals survive to the next generation.

<u>Uniform Stochastic:</u> In this scheme, individuals are selected with uniform probabilities.

6.3.3 Variation Operators

Mutation: Mutation is a genetic operator by which small modifications are made to a single genotype or phenotype of a selected individual. For real-valued phenotype approaches, the mutation operator is usually of the form of a small Gaussian change to the phenotype. For binary coded genotypes, a mutation is a random bit flip.

Crossover: Crossover is where two or more individuals exchange information to create one or more new individuals. In its simplest form, a single random locus is selected as a slice point and the segments are exchanged between two individuals at the locus. Figure 2.10 shows an example of a one-point crossover. In contrast to the one-point crossover, a multi-point crossover is when more than one locus is selected and information is exchanged in segments between the loci points.

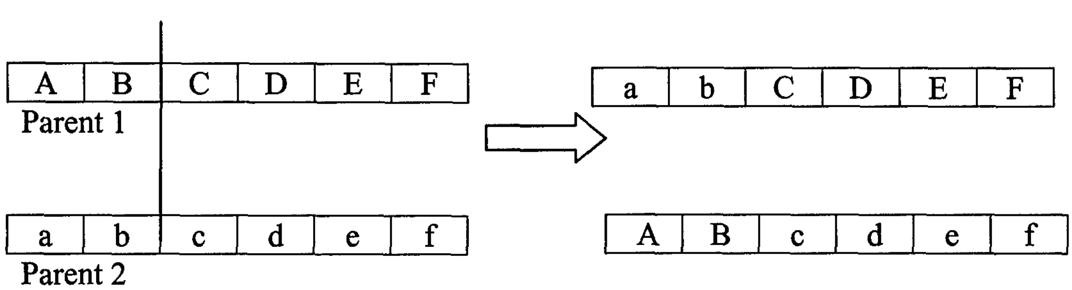

Figure 3: One-Point Crossover

6.4. ECAD-EA Methodology

The previous two sections described the main tools i.e., TIN and EA, used in this paper to identify an effective course of action in a dynamic uncertain situation. As discussed earlier, the term course of action in the context of this paper means the identification of a set of actions to be selected in a plan, their sequences, and their time of executions that maximize/minimize a certain metric (or set of metrics) related to the desired effect(s). The proposed approach that accomplishes this task is termed as ECAD-EA (Effective Courses of Action Determination Using Evolutionary Algorithms). Besides a completely specified TIN, the approach requires a system modeler to specify several other items before it determines an effective course of action. Some of these items are given below followed by their explanations, while Figure 4 shows the inputs and outputs of the ECAD-EA methodology in the form of a block diagram.

(a) Windows of Opportunity and Observation
(a) Identification of a Robust Metric
(b) Single or Multiple Desired Effects
(c) Constraints among Actionable Events

6.4.1 Windows of Opportunity and Observation

The proposed ECAD-EA methodology requires a system modeler to specify two parameters that aid in creating a boundary around the solution space. The first parameter is the duration in which all the actions must take place. This duration is referred to as either *window of opportunity* or *action window*. It should be mentioned that the window of opportunity refers to the time interval during which the whole plan should be executed. If some actions have more strict temporal requirements for their execution, then these requirements can be specified with the help of the constraint specification language discussed in subsection 7.4.5. The second parameter is the time period in which a decision maker is interested in getting the desired results. This duration is referred to as *observation window*. The two terms are explained further towards the end of this section with the help of an example.

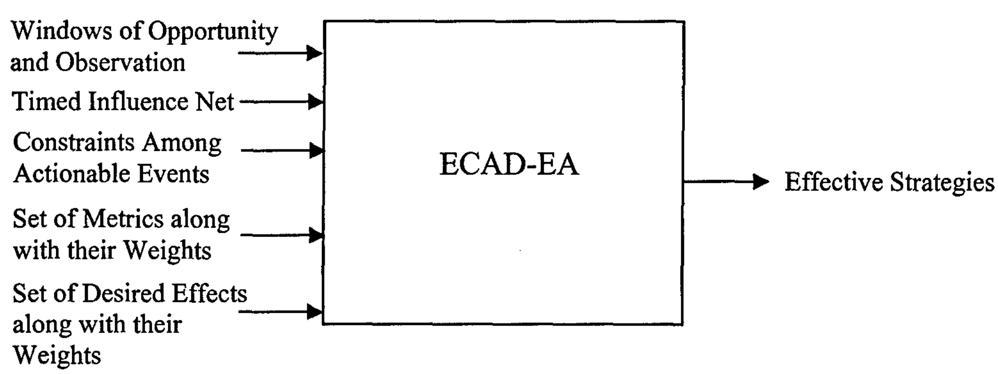

Figure 4: Block Diagram of ECAD-EA

6.4.2 Identification of a Robust Metric

To judge the fitness of a COA, a set of metrics is needed that can be used as Measures of Performance (MOPs). All the metrics discussed in this paper evaluate a COA in terms of the desired effect's probability profile produced by that particular COA. The fitness of a COA can be measured using either a simple or a composite metric. In case of a composite metric, weights should be assigned to each of the individual metric. Weights assigned to a metric may vary from situation to situation and is upon the discretion of the system modeler. A selected metric (or set of metrics) guides the EA through probability profiles' space while the EA searches for an effective COA. There are several features of a probability profile that can be considered during its evaluation. A preliminary list of features have been identified that could constitute potential metrics. The list includes the following and is explained with the help of probability profile of Figure 5.

1) <u>Maximum Probability Achieved:</u> This metric specify the highest probability achieved by the probability profile of a desired effect during the window of observation. For instance, in the profiles of Figure 5 the highest probability produced by COA #1 is 0.9 in contrast to 0.8 produced by COA #2.

2) <u>Time to Reach the Maximum Probability:</u> This metric is related to the first one and it specifies the time at which the maximum probability is achieved. For instance, COA #1 in Figure 5 has produced a better probability value (0.9) but the value is achieved at time 8. COA #2, on the other hand reaches the probability of 0.8 at time 6. If Metric 1 is considered alone then COA #1 would be considered better, but if a decision maker is interested in getting the probability of desired effect above a certain threshold, then COA #2 might be considered better if it satisfies the threshold requirement.

3) <u>Probabilities at Specific Points/Intervals:</u> A decision maker might be interested in profiles that have probabilities above (below) a certain threshold or highest at specific time instances or intervals. For instance, in Figure 5, a decision maker might be interested in those profiles which have higher probabilities during the time interval 4 to 8. COA # 1, in this case, is superior to COA #2. On the other hand, if the interval under consideration is between time 8 and 11, then COA #2 would be considered better. The time instances/intervals could be a list of values, i.e., the decision maker requires the probability to be above (below) certain thresholds at time 3, 9, and 13.

4) <u>Area Under the Curve (AUC):</u> This metric represents the area under a probability profile. The computation of AUC for probability profiles of Figure 5 is shown in Table 2. Based on the metric, it can be said that COA #2 is better than COA #1.

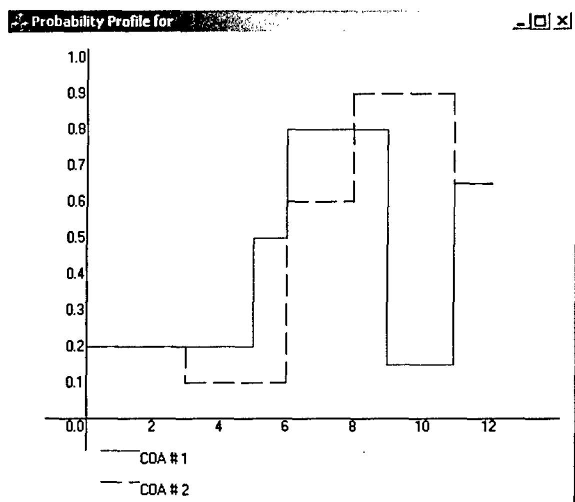

Figure 5: Two Competing Courses of Actions

Once a set of metric is specified, along with their corresponding weights, the fitness of a particular COA is measured as a weighted sum of the selected metrics. In other words,

$$\text{fitness(COA x)} = w_1\, m_1 + w_2\, m_2 + \ldots\ldots\ldots\ldots\; w_n\, m_n \qquad (1)$$

where w_1 is the weight corresponding to the metric m_1. In such cases when multiple metrics are considered, each metric is normalized on the scale of 0 to 1.

6.4.3 Single/Multiple Desired Effects (or Objectives)

A given problem could have either single desired effect or multiple (at times conflicting) desired effects. In case of multiple desired effects, a system modeler needs to weight the importance of each desired effect. The objective function then becomes a weighted average certain functions of each desired effect's probability profile. It should be noted that each profile can be evaluated using a different set of metrics and their corresponding weights. For instance, if there are two desired effects d1 and d2, then a decision maker might be interested in using Metric 4 (AUC) for d1 while using Metric 1 and 2 for d2.

Table2: Computation of AUC Metric for COAs in Figure 3.16(a)

COA # 1				COA # 2			
Time at which Change Occurs	Duration of the Interval **d**	Probability Of the Target Node **t**	Area **d x t**	Time at which Change Occurs	Duration of the Interval **d**	Probability Of the Target Node **t**	Area **d x t**
5	5	0.2	1	3	3	0.2	0.6
6	1	0.5	0.5	6	3	0.1	0.3
9	3	0.8	2.4	8	2	0.6	1.2
11	2	0.15	0.3	11	3	0.9	2.7
Sum			**4.2**				**4.8**

Equation 1 can thus be modified as

$$\begin{aligned}\text{fitness(COA x)} = & \; w_1 (w_{11}\, m_1 + w_{12}\, m_2 + \ldots\ldots\ldots + w_{1n}\, m_n) \\ & + w_2 (w_{21}\, m_1 + w_{22}\, m_2 + \ldots\ldots\ldots + w_{2n}\, m_n) \\ & \vdots \\ & + w_k (w_{k1}\, m_1 + w_{k2}\, m_2 + \ldots\ldots\ldots + w_{kn}\, m_n) \end{aligned} \tag{2}$$

when there are k number of desired effects and each effect can be evaluated using n number of metrics. w_i represents the weight assigned to the desired effect i, while w_{ij} represents the weights assigned to metric j for the same effect.

6.4.4 Constraints among Actionable Events

Several temporal and logical constraints might exist among the actionable events that are not directly modeled using TIN. The presence of these constraints results in the need of a constraint specification language that helps a system modeler in modeling temporal/logical constraints. The paper proposes one such language that assists a modeler in specifying these types of constraints. The consistency of these constraints can be checked by propositional and temporal logic based systems depending upon the nature of the given constraints. Table 3 describes the constructs of the proposed constraint specification language (CSL). Constraints 1 and 2 are for the cases where it is required that some actions must be in a specific state. Constraint 3 requires that A1 and A2 should have the same state (whether "True" or "False"). Constraint 4 specifies the time 't' at which action A must take place while, Constraint 5 shows that action A1 must happen before action A2. If the value of *v* is zero then A1 should happen earlier than A2 but if *v* is non-zero (positive) then A1 must happen at least *v* time units before A2. Constraint 6 specifies that two actions A1 and A2 should happen at the same time.

Table 3: Constructs of the Constraint Specification Language

1. True	A			// Factual Constraint
2. False	A			// Factual Constraint
3. Same_State	A1	A2		// Causal Constraint
4. At_Time	A	t		// Temporal or Factual Constraint
5. Before	A1	A2	*v*	// Temporal or Causal Constraint
6. Equal_Time	A1	A2		// Temporal or Conceptual Constraint

6.4.5 Specifics of the Evolutionary Algorithm

Once all the inputs to the ECAD-EA methodology are specified, the proposed approach searches the probability profile solution space to determine effective strategies. Consider the TIN shown in Figure 6. There are two actionable events, A and B, and one desired effect H. Suppose a decision maker is interested in selecting a course of action that maximizes the probability of H over a given time interval. The window of opportunity is set as an interval between time instances 1 and 10 while the observation window is set between time instances 1 and 20. Thus, the solution space consists of $2^2 \times 10^2 = 400$

probability profiles. Figure 7 shows the solution space when both actions A and B are true given that the fitness of a profile is measured using AUC metric. As it can be seen from the figure, the global maximum lies in the region when action B is taken earlier than action A.

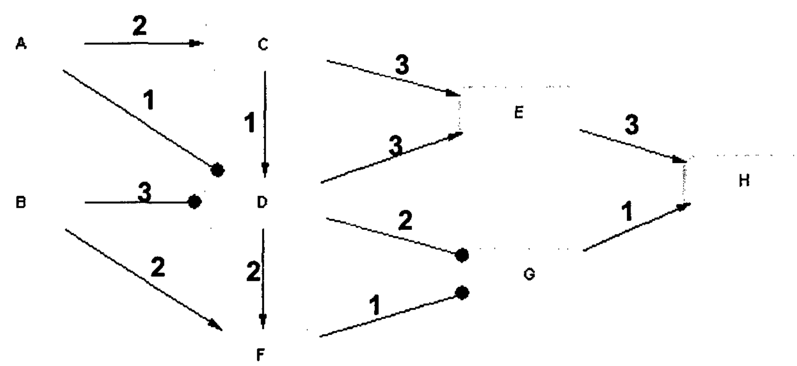

Figure 6: A Sample Timed Influence Nets

The proposed methodology explores the solution space of Figure 7 using an Evolutionary Algorithm. As described in Section 7.3, the main characteristics of an EA are the representation of individuals in the population, the selection mechanism, and the variation operators. The population in the EA of ECAD-EA consists of candidate courses of action that maximize a metric. Each individual in the population consists of actions and the times at which the actions are executed. Due to the fact that all the variables in a TIN have binary states, a bit representation is used for the actionable events. An integer value representation is used for the time at which an action is executed. Thus an individual (phenome) in the population of solutions has the following structure:

<Action 1, time (Action 1), Action 2, time (Action 2), ..., Action n, time (Action n)>

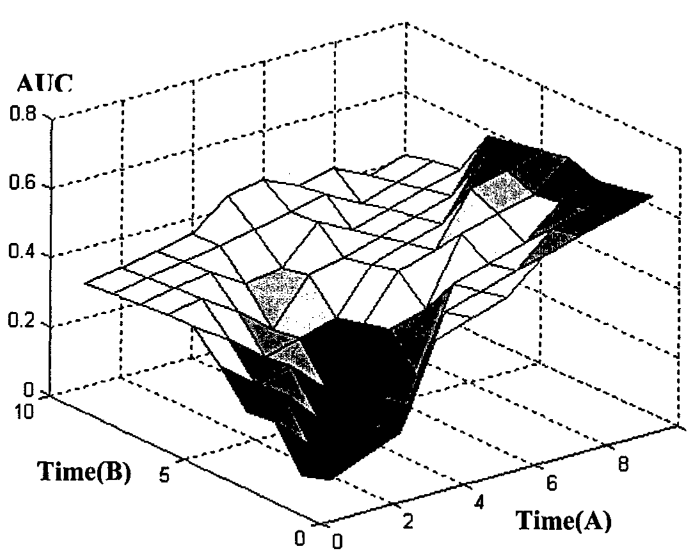

Figure 7: The Solution Space of the TIN Shown in Figure 6

For the example under consideration, the AUC metric is used to evaluate the fitness of a probability profile. The parent selection scheme is uniform stochastic, while the survival selection scheme is binary tournament. The last major ingredient of the EA is the variation operators. Since all the actions have binary states, a bit flip mutation operator is used for them, while a delta x mutation operator is used for the times at which actions are executed. The size of the population is kept at 10, while the EA is run till 20 generations. The details of the EA are summarized in Table 4.

Table 4: Parameters of the Evolutionary Algorithm

Representation: Phenotype
Parent Selection: Uniform Stochastic
Survival Selection: Binary Tournament
Mutation: Bitflip / Delta x
Mutation Rate: 0.2
Crossover: one point
Population Size: 10
Number of Generations: 20

The EA of Table 4 is applied on the TIN of Figure 6 along with the constraints regarding windows of opportunity and observation to find courses of action that maximize the AUC metric for event H's probability profile. Several solutions with very close fitness values are produced. Out of those solutions, the four top ones are shown in Table 5. The solution shown in the first row is read as follows: "in order to maximize AUC, action A is made True at time 9, while action B is made True at time 5". The other rows in the table can be read in a similar manner. The results agree with the intuition developed earlier by looking at the solution space of Figure 7. All the solutions are centered on the region where the global maximum lies. The solutions also conformed to the temporal relationship exhibits by Figure 7. In all the solutions, action B is executed before action A. Thus, the advantage of having multiple solutions, having very close fitness, is that the solutions can be generalized to produce a plan that describes time at a qualitative level or at a more general quantitative level. A rather detailed discussion on this issue is presented in Section 7.4.6.

Table 5: Four Best Solutions for TIN of Figure 6

A	time(A)	B	time(B)
True	9	True	5
True	9	True	6
True	8	True	5
True	9	True	4

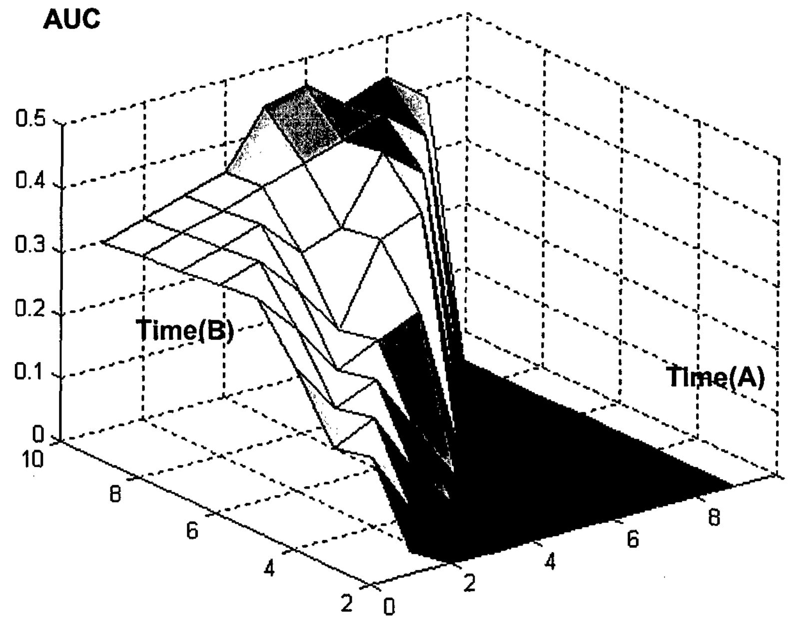

Figure 8: The Constrained Solution Space of the Model Shown in Figure 6

The solutions of Table 5 are produced without considering any user-defined constraints. Suppose it is required that a feasible solution must satisfy the constraint that A must be executed before B. Using the constructs of the proposed CSL, the constraint can be written as:

Before A B 0

The reduced solution space (based on AUC metric) is shown in Figure 8. When the EA of Table 4 is run on this modified solution space, it generates courses of action that maximize the AUC metric for event H's probability profile under the given constraint. The four top courses of action are shown in Table 6. Like the previous case, the solutions produced by the EA are in the region that is close to the global maximum of Figure 8.

Table 6: Four Best Solutions for the TIN of Figure 6 with Constrained Solution Space

A	time(A)	B	time(B)
True	8	True	10
True	7	True	9
True	6	True	9
True	4	True	9

6.4.6 Generalization of Temporal Relationships

The purpose of generating alternate courses of action is to determine any pattern that exists among the alternate courses of action. If there exists a pattern among the solutions that can be generalized, then this information can be very helpful to a decision maker as it gives more flexibility to him while he plans a course of action. This section presents a scheme for generalizing the temporal relationships that exist among the actions in the solutions, produced by the EA, having very close fitness value. Consider the three courses of actions shown in Table 7. There are four actionable events, namely, A, B, C, and D. It should be noted that actions in a TIN are instantaneous. Thus, only two kinds of temporal relationships can exist between actionable events, namely, 'Before' and 'Equal'. Thus, a point in time can be either 'Before' another point in time or it can be 'Equal' to some other point. For a detailed discussion of possible temporal relationships between points and interval, refer to [Zaidi, 1999;Zaidi and Levis, 2001].

Table 7: Three Alternate Courses of Actions

A	time(A)	B	time(B)	C	time(C)	D	time(D)
T	1	F	14	T	8	T	5
T	8	F	14	T	12	T	13
T	4	F	5	T	13	F	3

The temporal relationship that exists among the actionable events can be represented in the form of a grid, as shown in Table 8. In each cell, there are 3 entries that correspond to the 3 solutions presented in Table 7. For instance, in all 3 solutions time(A) is always less than time(B). Thus, the relationship between A and B in all cases is "Before" as shown in the cell corresponding the intersection of A and B (1st Row and 2nd Column). Similarly, the relationship between A and C is "Before" in all three cases as time(A) is less than time(C) in all solutions. D, on the other hand, is "Before" A in the last solution while A is "Before" D in the first two solutions. Other entries in the table can be read in a similar fashion. From the table, it can be seen that there are some cells where corresponding actions have a similar type of temporal relationship in all the solutions. Thus, the temporal relationship between those actions can be generalized. For instance, the following generalization can be depicted from Table 8.

A Before B

A Before C

D Before B

The above information aids a decision maker during the planning stage of a mission. It simply states that any plan that is developed to accomplish a particular objective should satisfy the above mentioned temporal constraints. Figure 9 shows an instance of possible relationships among the four actionable events on a single time line. The generalization of temporal relationships can be extended further. For instance, in the temporal relationship of Table 7, whenever D is in state True, it happens after A. On the other hand, when D is in state False, it happens before A. Thus, the temporal relationship between two actions can be generalized using If-Then type of rules. For instance, the following rules can describe the temporal relationship that exists between A and D.

Table 8: Temporal Relationships among Actionable Variables

	A	B	C	D
A		Before Before Before	Before Before Before	Before Before --
B			-- -- Before	
C		Before Before --		-- Before --
D	-- -- Before	Before Before Before	Before -- Before	

If D is True Then

A is Before D

If D is False Then

D is Before A

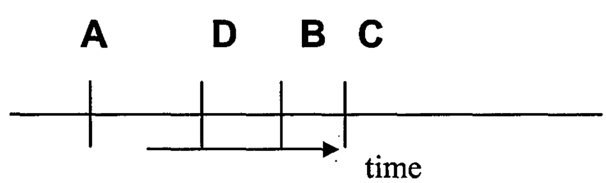

Figure 9: Actionable Events on a Single Time Line

6.5 Application of the ECAD-EA Methodology

This section applies the ECAD-EA methodology on a real model. The TIN model, shown in Figure 10, is developed to capture the events and the associated uncertainties in combating insider threat[20]. It should be mentioned that only the structure of the TIN is shown in the figure. The quantitative strengths and the time delays associated with each arc are not shown to keep things simple. The reader however, should not have any difficulty in understanding the results presented in this section. The purpose of building the insider threat model is twofold: to identify and analyze the actions that could be used by an insider to become a security risk for an organization and to analyze the actions taken by the organization to prevent any potential damage caused by the insider. In other words, a system modeler is interested in identifying an effective strategy that maximizes

[20] The model was presented to OSD-C3I on Dec 19, 2003 by Lee Wagenhals and Larry Wentz.

the likelihood of achieving the desired effect: "Insider Does Not Have Strength of Resolve to Attack". There are ten actionable events, drawn as the root nodes, in the insider threat model:

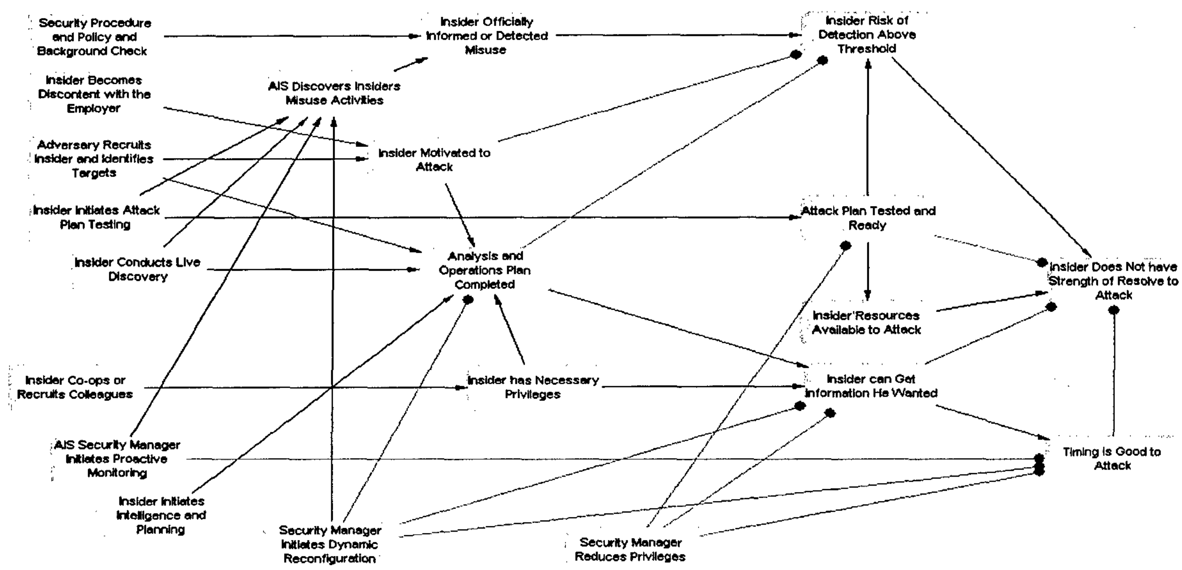

Figure 10: Influence Net for Insider Threat Model

Security Manager Reduce Privileges	(A)
Security Manager Initiates Dynamic Reconfiguration	(B)
Insider Initiates Attack Plan Testing	(C)
AIS Security Manager Initiates Proactive Monitoring	(D)
Insider Co-ops or Recruits Colleague	(E)
Insider Conducts Live Discovery	(F)
Insider Initiates Intelligence and Planning	(G)
Adversary Recruits Insider and Identifies Target	(H)
Insider Becomes Discontent with Employer	(I)
Conduct Security Procedure and Policy and Background Check	(J)

The labels next to the actions are used in the sequel to refer to the corresponding actions. The EA, as explained in Section 4.5, is applied on the TIN developed for insider threat. The parent selection scheme is set to uniform stochastic, while the survival selection is set to truncation. Population and generation sizes are set to 20 and 15, respectively. The window of opportunity is between 1 to 20 time units, while the window of observation is between 1 to 25 time units. The top 4 results produced by the ECAD-EA methodology are presented in Table 9. There are 20 elements in each solution. The odd elements (1, 3, 5, ...) represent the state of the actionable events, while the even elements (2,4,6,....) represent the time stamps at which the state of the actions are decided. The order of the actions in a particular solution is same as described above. The solution described in the first row of Table 9 says that "Security Manager Reduce Privileges" is made True at time

9 while "Security Manager Initiates Dynamic Reconfiguration" is made True at time 4 and so on. . Figures 11(a) and 11(b) show the comparison of Solution 1 to Solution 2 and Solution 1 to Solution 4, respectively. The probability profiles produced by all the solutions are close, i.e., it is hard to say which one is better than the other. In this kind of situation, the best thing is to produce a certain percentage of the top solutions (options) and then let the decision maker consider other characteristics which are not explicitly mentioned while planning a course of action in making his choice.

Table 9: Four Best Solutions for the TIN of Insider Threat Model

S.#	
1.	1 9 1 4 0 8 1 2 0 19 0 12 0 11 1 17 0 1 1 8
2.	1 6 1 1 1 11 1 8 0 6 1 9 1 12 0 5 1 14 1 6
3.	1 4 1 2 0 18 1 3 0 3 1 9 1 13 1 19 0 14 0 6
4.	1 1 1 7 0 4 1 12 0 6 0 11 0 19 0 15 0 16 0 9

The temporal relationship that exists among the solutions of Table 9 can be generalized to produce a feasible time line that aids a decision maker during the planning stage of a mission. Applying the same approach as explained in Section 4.6, the following temporal relationships are found among the 4 solutions shown in Table 9:

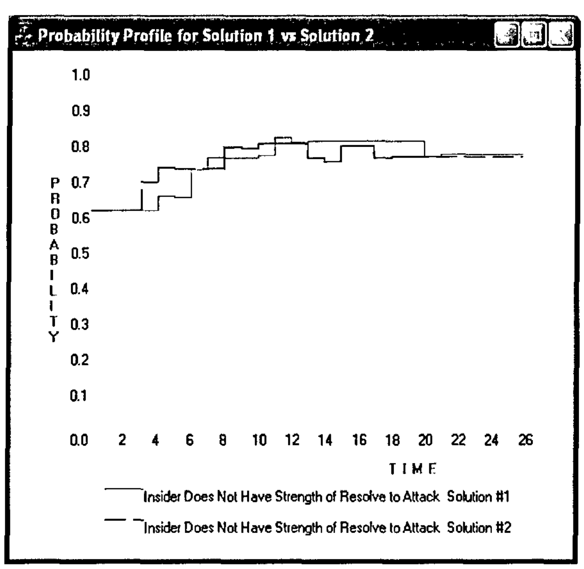

Figure 11: Comparison of Different Solutions: (a) Solution 1 vs. Solution 2

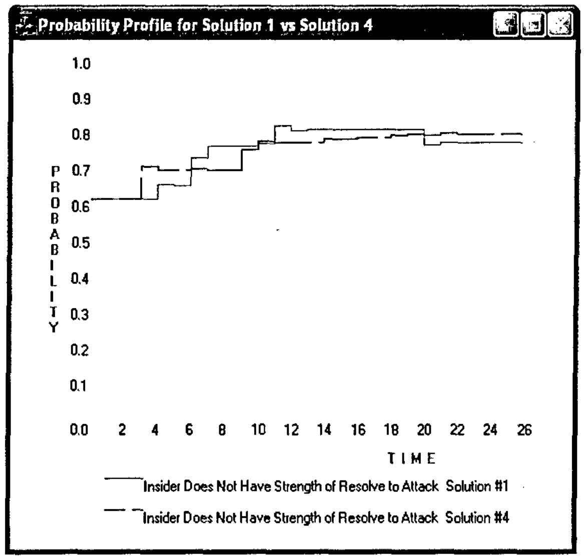

Figure 11: Comparison of Different Solutions: (b) Solution 1 vs. Solution 4

A Before F	A Before G	B Before A
B Before G	B Before H	B Before J
D Before G	J Before F	J Before G

A possible time line based on the above temporal relationships is shown in Figure 12. The above generalization also matches with the intuition one could develop after looking at the set of actions present in the problem. The temporal relationships roughly say that the actions taken by an organization (A, B, D, and J) should be taken before the actions taken by an adversary (F, G, and H).

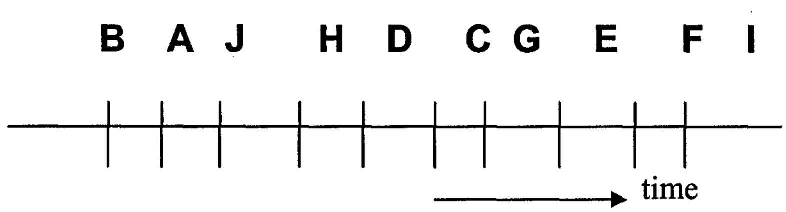

Figure 12: Actionable Events on a Single Time Line

It can be seen from the probability profile of Figure 11(a) that Solution #2 is in fact better than Solution #1 during the time interval 2 and 11. However, during the whole duration of 25 time units, Solution #1 turns out to be better based on the AUC metric. If a system modeler is interested in only the probability profile during the time interval 2 and 11, then the solutions produced by the ECAD-EA would be different. To demonstrate this fact, the EA is run again on the insider threat model, but this time the fitness of a solution is evaluated by the area under a probability profile during time interval 3 and 12. The top 4 solutions are shown in Table 10, . The probability profiles produced by the two top solutions are shown in Figure 13(a). Again, there is not much difference between the performances of the two solutions. Figure 13(b) compares the profiles of the best solutions produced using AUC and Interval metric. As expected, the solution based on interval metric performs better during the time interval 3 and 12.

Table 10: Four Best Solutions based on Interval Metric for the Insider Threat Model

S.#	
1.	1 1 1 2 0 6 1 1 0 17 1 13 0 2 0 18 0 15 1 10
2.	1 11 1 1 1 10 1 1 0 4 0 1 1 17 1 19 0 19 1 10
3.	1 4 1 2 1 19 1 17 1 5 0 3 1 14 1 11 0 3 0 13
4.	1 2 1 4 0 12 0 5 0 11 1 4 0 6 1 16 0 15 1 17

Table 11: Four Best Solutions for Insider Threat Model Having Few Constraints

S.#	
1.	1 1 1 1 0 15 1 4 0 12 1 6 0 9 0 8 0 1 1 1
2.	1 1 1 1 1 15 1 4 0 12 1 6 0 9 0 8 0 5 1 1
3.	1 1 1 1 0 10 1 4 0 12 1 6 0 2 0 14 0 1 1 1
4.	1 1 1 1 0 15 1 4 0 12 1 6 1 8 0 8 0 5 1 3

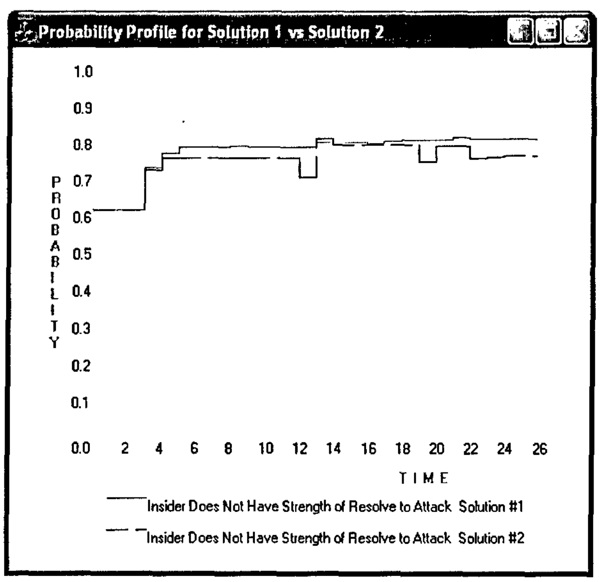

Figure 13: Comparison of Different Solutions: (a) Solution 1 vs. Solution 2

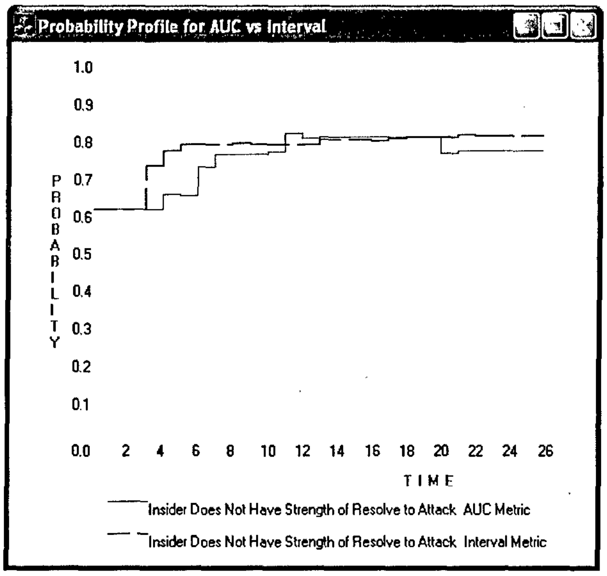

Figure 13: Comparison of Different Solutions: (b) Solution 1 vs. Solution 4

The solutions of Table 11 are produced without considering any constraints. Suppose it is required that 'Security Manager Initiates Dynamic Reconfiguration' must be true. Based on the temporal constraints, 'Insider Initiates Intel and Planning' must be at least 3 time units before 'Insider Initiates Attack Plan Testing'. During the plan execution, it turns out that 'Insider Conducts Live Discovery' at time 6. Furthermore, 'Security Manager Initiates Proactive Monitoring' at time 4. Finally, it is discovered at time 12 that the insider attempts to "Co-ops or Recruits Colleague" failed. All of these constraints can be written in terms of the Constraint Specification Language described earlier.

Before	'Insider Initiates Intel and Planning' 'Insider Initiates Attack Plan Testing'	3
True	'Security Manager Reduce Privileges'	
At_Time	'AIS Sec Mgr Initiates Proactive Monitoring'	4
True	'Insider Conducts Live Discovery'	
At_Time	'Insider Conducts Live Discovery'	6
False	'Insider Co-ops or Recruits Colleague'	
At_Time	'Insider Co-ops or Recruits Colleague'	12

The top four best solutions produced by the ECAD-EA methodology when applied on the insider threat model having the above constraints are presented in Table 11. The results presented so far consider only single desired effect, but the proposed approach works equally well on problems having multiple desired effects.

6.6 Conclusions

This section presents a methodology to identify effective courses of action in complex uncertain situations. The approach works on Timed Influence Nets that are used to model the uncertainties present in such complex situations. The presented approach uses Evolutionary Algorithms to identify effective courses of action. The approach not only provides a single best solution but it also provides several alternate solutions that are close enough to the best solution. These alternate solutions aid a decision maker in understanding the impact of actions' time of executions on the desired effect. The paper suggests a temporal generalization based on the similarities that exist in the temporal relationships among the actions.

A course of action is evaluated based on some pre-defined metrics. The paper discusses a set of metrics that can be considered, either individually or jointly (depending upon the situation), while evaluating a course of action. A constraint specification language that aids a system modeler in specifying the temporal and causal constraints among the actionable events is also presented.

References

Boyen, X. and Koller, D., "Tractable Inference for Complex Stochastic Processes," Proceedings of the Conference on Uncertainty in AI, 1998.

Chang, K. C., Lehner, P. E., Levis, A. H., Zaidi, S. A. K. and Zhao, X., "On Causal Influence Logic," George Mason University, Technical Report Dec. 1994.

Charniak, E., "Bayesian Network without Tears," in AI Magzine, vol. 12, 1991, pp. 50-63

Cooper, G. F., "The Computation Complexity of Probabilistic Inference Using Bayesian Belief Networks," Artificial Intelligence, vol. 42, pp. 393-405, 1990.

Dagum, P. and Luby, M., "Approximating Probabilistic Inference in Bayesian Belief Networks Is Np-Hard," Artificial Intelligence, vol. 60, pp. 141-153, 1993.

DeJong, K. A., "An Analysis of the Behavior of a Class of Genetic Adaptive Systems", Doctoral Dissertation," University of Michigan, 1975.

DeJong, K. A., Evolutionary Computing: MIT Press, 2005 (To Appear).

Fogel, L. J., Owens, A. J. and Walsh, M. J., Artificial Intelligence through Simulated Evolution. Chichester, UK: John Wiley Inc., 1966.

Goldberg, D. E., Genetic Algorithms in Search, Optimization, and Machine Learning. Massachusetts: Addison-Wesley, Reading, 1989.

Haider, S. and Levis, A. H., "A Heuristic Approach for Belief Updating in Timed Influence Nets," Proceedings of 2004 Command and Control Research and Technology Symposium, San Deigo, 2004.

Haider, S. and Zaidi, A. K., "Transforming Timed Influence Nets into Time Sliced Bayesian Networks," Proceedings of 2004 Command and Control Research and Technology Symposium, San Diego, 2004.

Haiying, T., Levchuk, Y. N. and Pattipati, K. R., "Robust Action Strategies to Induce Desired Effects," IEEE Transactions on Systems, Man, and Cybernatics - Part A: Systems and Humans, vol. 34, pp. 664-680, 2004.

Hanks, S., Madigan, D. and Gavrin, J., "Probabilistic Temporal Reasoning with Endogenous Change," in Proceedings of the Eleventh Annual Conference on Uncertainty in Artificial Intelligence (Uai-95), 1995, pp. 245--254.

Kjaerulff, U., "A Computational Scheme for Reasoning in Dynamic Probabilistic Networks," in Proceedings of the Eighth Conference on Uncertainty in Artificial Intelligence, 1992, pp. 121--129.

Neapolitan, R. E., Learning Bayesian Networks: Prentice Hall, 2003.

Pearl, J., Probabilistic Reasoning in Intelligent Systems: Network of Plausible Inference: Morgan Kaufmann, 1987.

Rosen, J. A. and Smith, W. L., "Influence Net Modeling with Causal Strengths: An Evolutionary Approach," Proceedings of the 1996 Command and Control Research and Technology Symposium, Monterey CA, 1996.

Santos, J. E. and Young, J. D., "Probabilistic Temporal Network: A Unified Framework for Reasoning with Time and Uncertainty," International Journal of Approximate Reasoning, vol. 2, 1999.

Schwefel, H. P., Evolution and Optimum Seeking: John Wiley & Sons, 1995.

Schwefel, H.-P., "Evolutionsstrategie und Numerische Optimierung," Department of Process Engineering, Technical University of Berlin, 1975.

Wagenhals, L. W., "Course of Action Development and Evaluation Using Discrete Event System Models of Influence Nets," PhD Dissertation, School of Information Technology and Engineering, George Mason University, 2000.

Zaidi, A. K., "On Temporal Logic Programming Using Petri Nets," IEEE Transactions on Systems, Man, Cybernetics A, vol. 29, 1999.

Zaidi, A. K. and Levis, A. H., "Temper: A Temporal Programmer for Time-Sensitive Control of Discrete Event System," IEEE Transactions on Systems, Man, Cybernetics A, vol. 31, 2001.

SECTION 7

Modeling for Future Command and Control Architectures

Holly A. H. Handley and Alexander H. Levis

7.1 Introduction

As the military moves to redesign command and control architectures to incorporate information technologies, models are necessary to predict the behaviors and performance of the proposed command structures. However, many of the performance metrics, such as speed of command and shared situational awareness, have not been included in previous command and control models. Models of command and control architectures have been developed over the last eight years in order to examine the behavior and performance of experimental command centers performing missions in a laboratory environment [Handley et al., 1999]. Each model met the requirements of the experiment and was validated with post experimental data; however, each model was limited to the conditions of the hypotheses and the performance metrics of the organizations and missions being developed.

A task process model has been developed as part of a recent pre-experimental modeling iteration for a subject experiment [Handley and Levis, 2003]. While previous models had focused on the decision maker process and metrics of the decision maker workload, the task process model captures the stages of a task over its lifetime, including information needs and decision maker activity required at each stage. By using the task process model, metrics for decision maker participation and information requirements regarding specific tasks can be elicited. The task process model was found to correlate well with tasks that required a single decision maker completing a task with a single resource. The model did not explain, however, why some tasks stop mid process and resume at a later time and it does not represent decision maker synchronization well, i.e., two or more decision makers supplying resources within a specified time window to complete a task. Both of these conditions require a decision maker model to work in conjunction with the task process model, in order to explicitly represent the interaction between the task and the decision maker.

In order to reconcile the task process model with an existing decision maker model, the five-stage interacting decision maker model [Levis, 1995], the empirical data collected from the subject experiment used to validate the task process model were examined. From the data, the task stages that required different decision makers were identified, along with delays in the task stages due to decision maker synchronization and interruptions in task processing due to engaged decision makers. The empirical data offered insights on how decision makers coordinated on complex tasks. An enhanced model was then created that combined the task process model and the decision maker model. The enhanced model will allow more sophisticated modeling of interactions between decision makers, such as decision maker synchronization and information sharing. By combining the task process model with the decision maker model, surrogate measures for speed of command and situation awareness can be developed and used to evaluate the behavior and performance of command and control information and decision processes, essential to assess any future command and control architecture.

The remainder of the paper is organized as follows: the next section describes the task process model while section 3.0 identifies its limitations. Section 4.0 describes the five stage decision maker model and section 5.0 describes the enhanced model that results from joining these two component models. Section 6.0 describes the performance measures used with this model, specifically speed of command and shared situational awareness; section 7.0 concludes the paper.

7.2 Description of the Task Process Model

The task process model was designed in conjunction with a subject experiment examining the relationship between different command and control architectures and alternative scenarios [Diedrich et al., 2002]. The task process model emulates the series of stages a task follows over its lifetime; each task that appears in the scenario is represented and evaluated separately. A task is an activity that entails the use of relevant resources and is carried out by an individual decision maker or a group of decision makers to accomplish the mission objectives or in defense of own assets. The task stages are based on the simulator used in the subject experiment, the Dynamic Distributed Decision-Making (DDD) simulator. The model was developed and validated using trial experimental data; see Handley & Levis [2003] for a complete description of the model development.

The task process model is shown in Figure 1. The first stage, Appear, occurs when a task (in this case a threat) is first present in the environment. This is controlled by the input scenario which specifies the time that each task appears. As soon as the task is noticed, either by a decision maker or a sensor, it is Detected and a decision maker initiates its processing. The task is then Identified; this indicates the decision maker knows what type of task it is and what type of resource (in this case a weapon) can be used to process the task. When the weapon is launched and travels to collide with the task, the task is defined as Attacked. When the resource has succeeded in completing the task (the weapon has destroyed the threat), the task is considered Destroyed. Lastly, the task Disappears from the simulator screen.

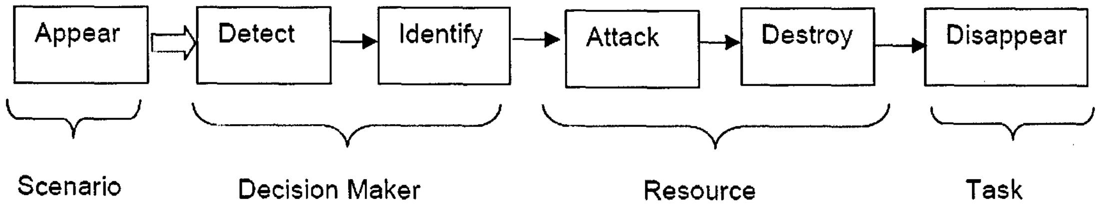

Figure 1: Task Process Model

The output of the task process model is a task completion time for every task in the scenario. Each stage of the sequential model has a delay determined by the attributes of the task, decision maker, or resource. Once the task enters the detect stage, it proceeds through the process uninterrupted, accumulating the delays at each stage until completion. The finish time of the task is the sum of the delays of each stage in the process; the task delay is the completion time minus the time the task first appeared, representing the actual processing time of the task:

$$t_{finish} = t_{appear} + t_{detect} + t_{identify} + t_{attack} + t_{destroy} + t_{disappear.} \quad [1]$$

$$t_{delay} = t_{finish} - t_{appear} \quad [2]$$

The time the task appears, t_{appear}, is predetermined by the scenario; the scenario is a list of all input tasks and the time they enter or appear in the model. The detect delay, t_{detect}, of each task is variable depending on the activity of the decision maker. If the decision maker responsible for the task is not currently acting on another task, the current task will be detected immediately. If, however, he is engaged with another task, the current task will wait until he is unoccupied; this variability is represented by the larger arrow in Figure 1. The delay associated with identifying the task, $t_{identify}$, represents the processing time by the decision maker. At the end of this stage the decision maker has identified the type of threat the task represents and the appropriate type of resource to use against it. This stage has a fixed delay to represent the decision maker's processing time. This value was determined empirically by comparing model simulation data at increasing levels of decision maker delay to trial experimental output data. There is also a workload limit of one task imposed on the model at this stage. In the Attack and Destroy stages the parameters of the resource chosen by the decision maker to counteract the task provide the delay times; t_{attack} is the launch delay of the resource and $t_{destroy}$ is the travel time of the resource to the location of the task. The delay of the Disappear stage, $t_{disappear}$, represents the delay between the time the threat is attacked and the time it disappears from the display; this value is specific to each task class.

The model was implemented using Colored Petri nets, a graphical modeling language and a powerful modeling tool used to expose critical time dependencies, task concurrencies and behavior that is event driven. The model was implemented in Design/CPN and simulated under different conditions as determined by the experimental design. While the DDD simulator creates the environment for the subject experiments, it also captures the subject's actions and task data throughout the course of the experimental scenario. This information is made available after the experiment in log files, which can be sorted by decision maker, resource, or task identifiers to find timing information. The model was validated after the subject experiment by comparing the timing of the task stages recorded in the log files with the timing of the task stages used in the simulation model.

Certain delays in the process are fixed based on the task type or the resource chosen while other delays are variable depending on the activity level of the responsible decision maker. The task pattern in the scenario elicits different decision maker activity levels depending on the architecture; the architecture determines what resources a decision maker controls and what types or locations of tasks each is responsible for. This model was used to predict congruence between architectures and scenarios. Scenarios varied the arrival time, type, and location of tasks, which in turn changed the loading on decision makers and affected his choice of resource; congruence was evaluated as the ability of the different architectures to process the scenario tasks in a timely manner. The task process model allows the evaluation of individual task delays; looking at a single task process allows the correlation of the task delay, resource used and decision maker engaged with the experimental data Looking across multiple task processes can be used to identify concurrency between tasks, decision maker workload at a particular time, and platform activity across time.

7.3 Limitations of the Single Task Model

Empirical data from the subject experiment [Handley and Levis, 2003] was used to validate the task process model. The performance of the model was validated by comparing the task completion times of the model to the experimental results. The average correlation between the model data and the experimental data was 0.86, with an average of 58 tasks per scenario (see Appendix A). While the final output of the task stages correlated well, there was a discrepancy between the modeled tasks and some of the experimental results. Examination of the time data of the experimental task stages indicated that often tasks were interrupted in the middle of the process and then were resumed later on. In order to verify the sequence of stages that composed the task process, experimental data from two task classes (threats) were examined in detail: enemy patrol boats and enemy air attacks. Both of these tasks required one decision maker and one resource to complete. Extraction of the empirical stage delays showed that the individual tasks fell into two categories: those that had an interruption, or a large time delay, in their task process, and those that did not. Note that the DDD simulator is a real time simulation, i.e., one second of real time represents one second of simulation time. Examples of the simulation time at each task stage are shown in Figure 2.

	Arrive	Detect	Identify	Select	Attack	Destroy
Patrol Boat #218						
Team MA	370	371	372	430	434	439
Team MC - Interruption	370	371	372	584	588	593
Air Attack #406						
Team SA	1251	1254	1322	1340	1342	1347
Team SE - Interruption	1251	1254	1322	1380	1389	1394

Figure 2: Example Tasks with Task Process Interruptions

The empirical data suggests that in some instances there is a break in the processing between the Identify and Attack stages. The stage Select was introduced in the log files to indicate the continuation of a task after an interruption. This break represents the decision maker disengaging from the current task to attend to another, higher priority task, before returning to the original task. This requires including another stage in the task process model, the Select stage, which represents another variable delay depending on the activity level of the decision maker. This also implies the need for a coupling of a decision maker model with the task process model to allow for variations in the task processing due to the activity level (workload) of the decision maker.

Many of the tasks in the scenario required the interaction of one or more decision makers to combine their resource in order to execute the task. These tasks were not included in the model simulation, but were present in the experiment and in the empirical data collected for examination. An example of the interaction of the decision makers synchronizing their resources is shown in Figure 3:

Stage	Time	DM	Resource
04-Select	1726	2	
04-Select	1756	3	
11-Attack	1759	2	SOF-500
12-Assist	1761	3	SOF-501
18-Destroy	1789	2	

Figure 3: Decision Maker Synchronization on Task 28, Team D Run S2

These synchronized tasks could not be included in the original task process model design; however, they can be addressed if the decision maker model is included explicitly.

7.4 A Five Stage Decision Maker Model

In order to study the behavior of an organization, it is necessary to have a model of its components, namely the individual decision makers. March and Simon [1958] hypothesized that decision makers follow a two step process: first determining the situation and then determining a response. This led to a two stage decision maker model by Wohl [1979] which was expanded to four stages by Boettcher and Levis [1982] in order to accommodate interactions between decision makers. Remy and Levis [1986] formalized these interactions. Levis [1992] presented a model of a five stage interacting decision maker that subsumed the previous models. This model presupposes that the decision makers are executing well-defined tasks for which they have been trained and that there is a limit to the amount of processing a decision maker can perform [Boettcher and Levis, 1982] in accordance with the bounded rationality constraint [March, 1978].

The five stage decision maker model is shown in Figure 4. The decision maker receives a signal, x, from the external environment or from another decision maker. The situation assessment stage (SA) represents the processing of the incoming signal to obtain the assessed situation, z, which may be shared with other decision makers. The decision maker can also receive a signal z' from another decision maker; z' and z are then fused together in the information fusion (IF) stage to produce z''. The fused information is then processed at the task processing (TP) stage to produce v. A command or control information from another decision maker is received as v'. The command interpretation (CI) stage then combines v and v' to produce the variable w, which is input to the response selection (RS) stage. The RS stage then produces the output y to the environment, and/or the output y' to other decision makers.

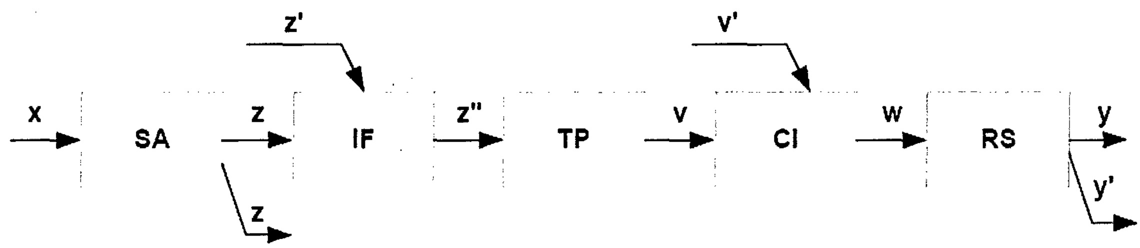

Figure 4: Five Stage Interacting Decision Maker

The model depicts explicitly the stages at which a decision maker can interact with other decision makers or the environment. A decision maker can receive inputs from the external environment only at the SA stage. However, this input x can also be from another decision maker (the y' output) from within the organization. A decision maker can share his assessed input through the z output at this stage. The z' input to the IF stage is used when the decision maker is receiving a second data input. This input must be generated from within the organization and can be the output of another decision maker's SA or RS stage. The fused information from the IF stage, z", is the input to the TP stage. The decision maker's function is performed at this stage and results in the output v. In the CI stage, the decision maker can receive control information as the input v'. This is also internally generated and must originate from another decision maker's RS stage. In the RS stage, an output is produced; y is the output to the environment and y' is the output to another decision maker. Thus the interactions between two decision makers are limited by the constraints enumerated above: the output from the SA stage, z, can only be an input to another decision-maker's IF stage as z', and an internal output from the RS stage, y', can only be input to another decision maker's SA stage as x, IF stage as z', or CI stage as v'.

A decision maker need not exercise all five stages when performing a task. Depending on the inputs and outputs required, a decision maker can instantiate different subsets of the five stage model.

7.5 Enhanced Task Process Model

The two limitations identified in the current task process model are the inability to allow a decision maker to disengage from a task in order to initiate the processing of another task and the inability to represent complex tasks, i.e., tasks requiring multiple decision makers to synchronize resources to accomplish a task. Both of these limitations require a coupling of the task process model with the five stage interacting decision maker model.

7.5.1 Task Interruption

Currently in the single task process model, the delays of the Detect and Identify stages are due to the decision maker; the decision maker is implicitly associated with these task stages. This relationship can be made explicit by associating the Detect-Identify stages of the task process model with the Situation Assessment (SA) – Response Selection (RS) stages of the decision maker model. The task process model is now constrained by the decision maker model during these stages. An additional task stage was identified in the empirical data: the Select stage preceded the Attack stage and was used to indicate that a decision maker had continued processing the interrupted task. The decision maker SA-RS stages can again be associated with the task process Select – Attack stages. This coupling of the task process model with the decision maker model is shown in Figure 5.

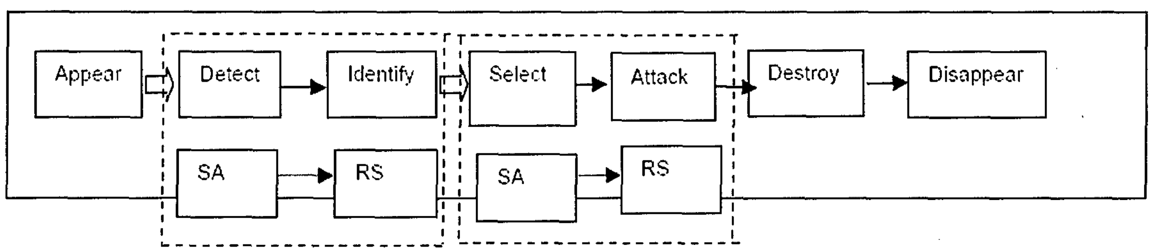

Figure 5: Single Task Model Coupled with Decision Maker Model

This coupling allows a second variability in the task processing: the original interruption in processing occurs between the Appear and Detect stages, which represents the variable delay due to an engaged decision maker, and an additional interruption between the Identify and Select stages, when a decision maker may disengage from the task in order to attend to another task. The Select stage delay is also a variable delay depending on the activity of the decision maker. The linear delay equations are modified by adding the additional variable delay term t_{select}:

$$t_{finish}' = t_{appear} + t_{detect} + t_{identify} + t_{select} + t_{attack} + t_{destroy} + t_{disappear.} \quad [3]$$

$$t_{delay}' = t_{finish}' - t_{appear} \quad [4]$$

7.5.2 Decision Maker Synchronization

The task process model is limited in that it can currently process only tasks requiring a single decision maker with a single resource. While these tasks did account for the majority of the tasks in the experimental scenarios, these were mostly tasks that defended own assets. The tasks that defined the objective of the scenario mission were the complex tasks that required two or three decision makers to synchronize their resources to accomplish the task within a defined time window; if any one decision maker applied his resource outside of that time window the task would fail.

In order to represent the synchronization of decision makers in the model, the different roles a decision maker assumes in processing different types of tasks must be defined in terms of the five stage model. Three decision maker roles have previously been identified [Levis et al., 1998]. The Independent role is defined as a decision maker acting on a task that he can then execute without interacting with other decision makers; this is the role of the single task in the current task process model. The Leader role is defined when a decision maker has to execute a task by interacting with other decision makers, however this decision maker is the initiator and sends synchronization messages to the other decision makers. The Follower role is defined for decision makers who must provide resources to execute a task with other decision makers, but it is another decision maker that sends the synchronization information.

The Independent role is used in the single task process model; a single decision maker with a single resource processes the task. Figure 4 and equations [3] and [4] describe this model. The Leader and Follower roles are bound together by the interactions required to synchronize their efforts. In all complex tasks, the commander's intent identifies a specific decision maker as the leader for that task. This decision maker will initiate the task and interact with the other decision makers as necessary to complete the task. These interactions are shown in Figure 6.

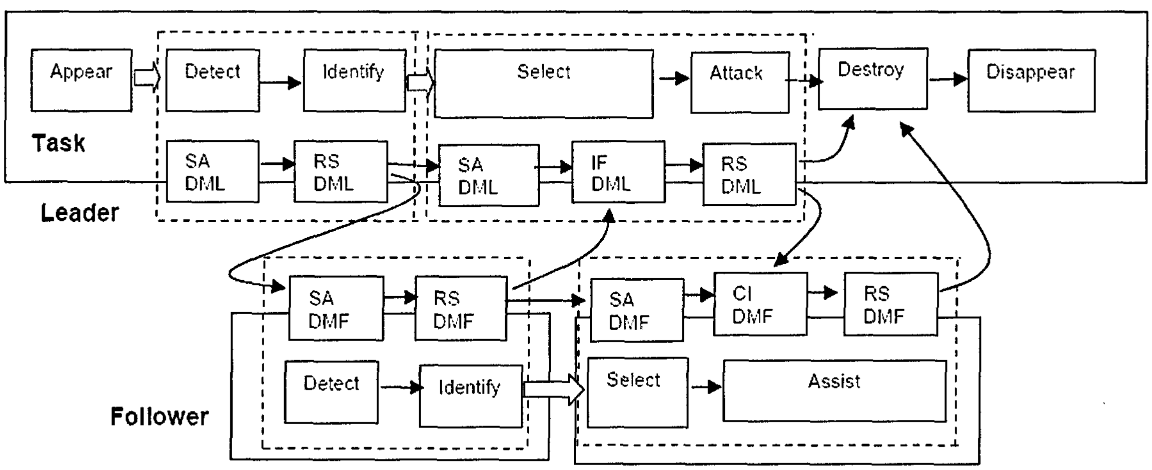

Figure 6: Leader and Follower Role Interactions

The Leader (DML) will first Detect and Identify the task; he identifies the additional resources and their responsible decision makers he needs for the complex task. The Leader is modeled with SA-RS stages, similar to the Independent role with the addition of the y\` output. The y\` output is used to alert the Follower decision maker(s) to Detect and Identify the task in preparation for a synchronized attack. The Follower (DMF) is modeled with the same SA-RS combination and the y\` output is used at the Leader's Select stage to indicate the readiness of the Follower decision makers; the Leader's Select stage is still variable dependent on the other tasks he is engaged with. The Leader will then begin the attack by launching his resource and signaling the other decision makers to synchronize their resource launch. The Leader's Select stage is modeled by a SA stage to indicate the task has been selected and an IF stage to include the "ready" signal from the Followers. The Attack stage is modeled as a RS stage, which indicates the launch of the resource and again a y` signal is sent to interact with the Follower. The Follower's Select stage also has a variable delay depending on his task priorities; the SA stage represents his Select stage. The Follower's task process stage is indicated as "Assist" in the DDD simulator; the Assist stage is modeled as a CI stage that waits for the synchronization signal from the Leader, and then a RS stage where the resource is launched.

The completion time of the task must be represented as a combination of delays from both the Leader and the Follower decision makers. In terms of the delays incurred by the Leader, the task completion time is:

$$t_{finish}`` = t_{appear} + t_{detectL} + t_{identifyL} + t_{selectL} + t_{IFL} + t_{attackL} + t_{destroy} + t_{disappear.} \quad [5]$$

However, t_{IFL}, which represents the delay associated with waiting for the "ready" response from the Follower decision maker is not a valid entry; it cannot be traced as a task stage in the DDD simulator. This delay can be represented, however, as the Detect and Identify stages of the Follower decision maker, $t_{detectF} + t_{identifyF}$; but these delays may be occurring concurrently with the Select delay of the Leader. The delay can then be correctly represented as a maximum of the pair:

$$t_{finish}`` = t_{appear} + t_{detectL} + t_{identifyL} + t_{selectL} + \max\,[0, (t_{detectF} + t_{identifyF}) - t_{selectL}] + t_{attackL} + t_{destroy} + t_{disappear.} \quad [6]$$

Likewise, the Destroy stage of the task is dependent on all synchronized resources arriving with the window of attack determined by the task class. In this case:

$$t_{assistF} - t_{attackL} < \Delta\, t_{attack} \quad [7]$$

where $\Delta\, t_{attack}$ is the pre-determined time completion window of the task; if this condition is not met the threat is not destroyed. For completeness:

$$t_{delay}`` = t_{finish}`` - t_{appear}. \quad [8]$$

The enhanced task process model was implemented as a Colored Petri net and re-simulated with the experimental scenario. The model output was again correlated with the experimental data; the average correlation over 12 teams was 0.80 with an average of 71 correlated tasks (see Appendix A). In order to add complex tasks to the model, the scenario was modified to accept the mission tasks. The mission tasks are complex tasks that must be completed in precedence order and require multiple resources; in some architectures one decision maker may own the complete set of resources, in other architectures he must coordinate with other decision makers to complete the set of resources required. While the existing scenario, the independent tasks, was an input list of tasks and their (fixed) arrival time, now the mission tasks are triggered by the completion of other tasks and as such have no fixed arrival time. This makes the correlation more complex as not only is the completion time variable, but also the actual appear time.

The delay times of this model are dependent on the interaction of the architecture and the scenario, and on the interaction of the Leader and Follower decision makers in complex tasks. The decision maker SA delay times (t_{detect}, t_{select}) represent the delay for the decision maker to commence work on the task, either initially or after an interruption. This delay depends on what other tasks the decision maker is engaged on and task priorities. The IF and CI stages are junctions where the decision makers exchange information; task processing suspends until all information is exchanged. It would be difficult to determine these compound delays without executing the model in simulation mode.

7.6 Performance Measures

7.6.1 Speed of Command

Speed of Command is defined as the time from when a threat is detected until it is engaged. A surrogate measure for the speed of command in the enhanced task process model is the task delay, i.e. the difference between the completion time of the task and the time the task appeared; the time from the Detect stage to the Disappear stage in the model. This is equation [4] for a single task and equation [8] for a synchronized task. Tasks can be evaluated individually using this metric as in Figure 6, or the accumulated task delay over the course of the scenario can be compared across architectures, as in Figure 7.

The data in the graphs show the simulation results for two different architectures, termed Functional and Divisional; in both cases the same scenario was used. The delay of each task versus the time it appears for each architecture is shown in Figure 6. While this graph shows the differences in delay for each individual task, it does not give a good indication of how each architecture is performing with regards to Speed of Command. A better indicator of the performance is the accumulated delay of individual tasks as the scenario progresses over time as shown in Figure 7; the Divisional architecture shows an improvement of 17.8% in accumulated task delay, or Speed of Command, over the Functional architecture. The enhanced model includes the variability of the decision maker's attention to the task, not only the initial delay in the Detect stage, but also delays that may occur due to interruptions in the task processing at the Select stage. These results concur with the experimental output; for this scenario, termed the "M" scenario, the Divisional organization outperformed the Functional organization.

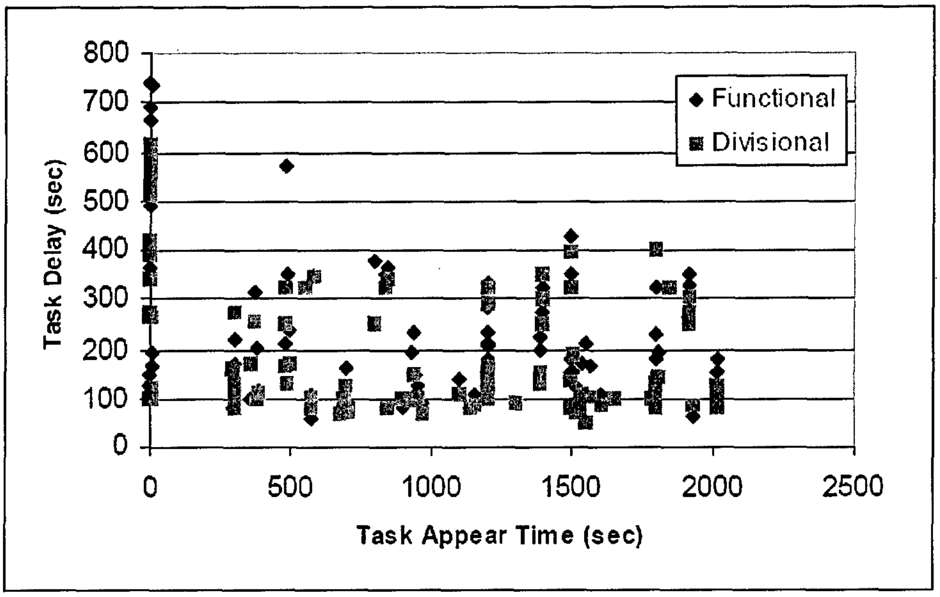

Figure 6: Speed of Command as Task Delay for Individual Tasks

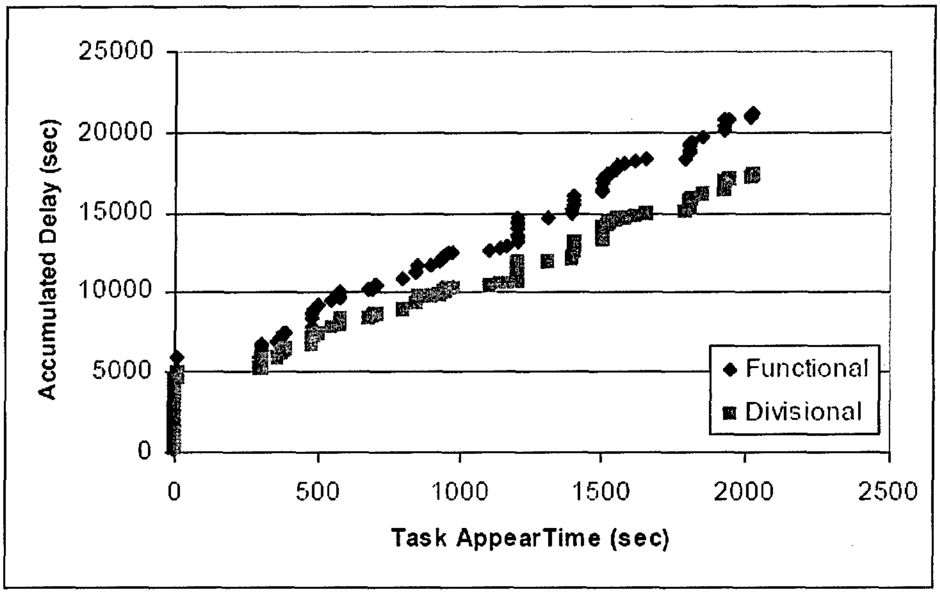

Figure 7: Speed of Command as Accumulated Task Delay for Scenario

7.6.2 Shared Situational Awareness

Shared Situational Awareness is the ability of a team of decision makers to perceive and understand a tactical picture that is complete and consistent across the team. In the single task process model there was no mechanism for decision makers to interact on a task; therefore there was no metric to gauge the situational awareness of multiple decision makers on the same task. The enhanced model specifically allows decision makers to synchronize their efforts to complete a task, which allows the opportunity to propose a metric to observe shared situation awareness.

On complex tasks that require multiple decision makers, a time window exists for each task in which all required resources must be fired. This allotted time can be described as a window of attack whose parameters are determined a priori by the requirements of the task; different task types may have different windows of attack. Two quantities are needed to specify the window of attack: the lower and the upper bounds of the time interval, t_s and t_f, respectively, or one of the bounds and the length of the interval, e.g. t_s and Δ t [Cothier and Levis, 1986]. The lower bound of the window is the time the first resource attacks the task and the length of the window is the predetermined time window of attack. In order for the attack to be successful, the time the final required resource attacks the task must be within the window's bounds:

$$t_f < t_s + \Delta t \qquad [9]$$

This window of attack (which is equivalent to [7]) can represents a surrogate measure for Shared Situation Awareness for the decision makers participating in the task. For the task to succeed the team of decision makers all need to apply the correct resources to the correct task within a finite period of time indicating a consistent and complete tactical picture. As the number of decision makers who participate in an attack increase, this metric becomes more meaningful.

The enhanced task process model can provide insights on Shared Situational Awareness based on the interactions between the Leader and Follower decision maker on synchronized tasks. There are three interaction points between the two roles: the output of the Leader's Identify (RS) stage to the Follower's Detect (SA) stage, the output of the Follower's Identify (RS) stage to the Leader's Select (IF) stage, and the output of the Leader's Attack (RS) stage to the Follower's Assist (CI) stage. While the first two will affect the overall delay of the task, the last interaction affects the final synchronization of the launch of resources. The critical point is the variable delay of the Follower's Select stage; if he delays too long resuming processing of the task, he will miss the window of attack initiated by the Leader's resource launch. This variable delay is a function of the architecture interacting with the scenario and can be used to compare across architecture scenario pairs; an example is shown in Figures 8 and 9.

Figure 8 shows both the Leader and Follower attack times for a set of 14 complex tasks, similar to those used in the "M" scenario above but in a separate, investigational scenario where the tempo of operations has been increased, completed by both Divisional (Div) and Functional (Fun) architectures. This can be used to evaluate how situational awareness varies over the course of this investigational scenario. In this case situational awareness seems to improve over the course of the scenario for both architectures. Figure 9 makes explicit the attack delay ($t_f - t_s$) versus the task window of attack (Δt_{attack}). The Functional architecture has seven tasks that miss the window and the Divisional has four.

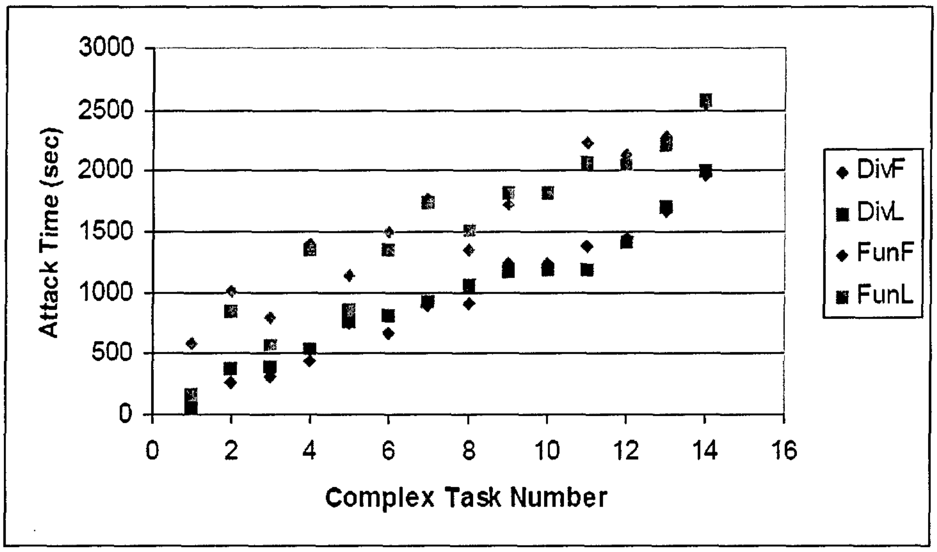

Figure 8: Shared Situational Awareness as Leader-Follower Attack Times

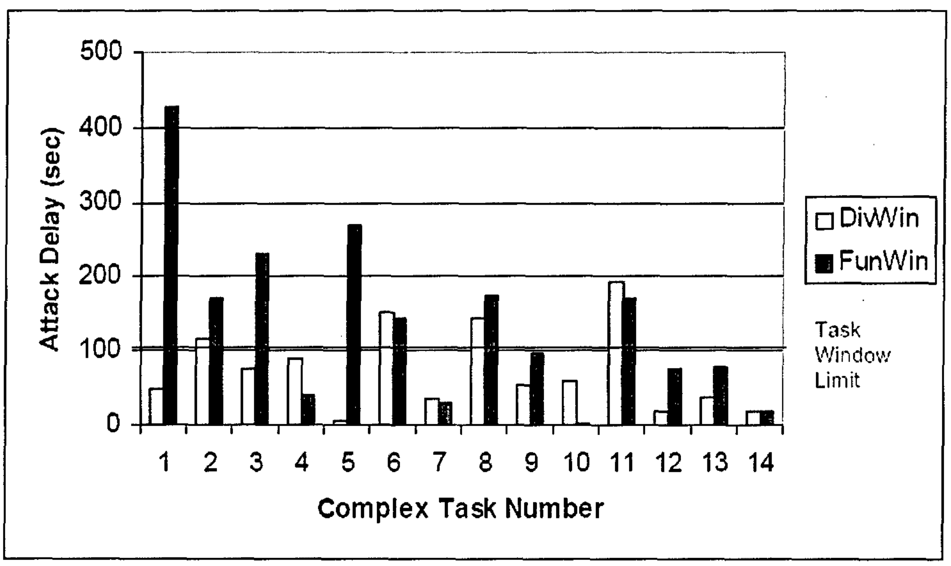

Figure 9: Shared Situational Awareness as Attack Delay

For a consistent Attack Window of 100 simulation seconds, the average window for the Divisional architecture is 73.8, while for the Functional architecture it is 136.4. In this case the Divisional architecture has a 46% improvement in Shared Situational Awareness. This metric was difficult to obtain from the subject experimental data, as many of the complex tasks were not attempted, and so no window data comparison was made.

7.7 Conclusion

The task process model was designed in conjunction with a subject experiment examining the relationship between different command and control architectures and alternative scenarios; the task process model emulates the series of stages a task follows over its lifetime. Limitations to the model were identified, specifically the inability to allow a decision maker to disengage from a task in order to initiate the processing of another task and the inability to represent complex tasks, i.e., tasks requiring multiple decision makers to synchronize resources to accomplish a task. Both of these limitations require a coupling of the task process model with a decision maker model. The five stage interacting decision maker model depicts the stages at which a decision maker can interact with other decision makers or the environment, i.e. the task process. The relationship between the models was made explicit by associating the Detect - Identify and Select - Attack stages of the task process model with the Situation Assessment (SA)-Response Selection (RS) stages of the decision maker model. The task process model is now constrained by the decision maker model during these stages.

The delay times of the enhanced task process model are dependent on the interaction of the architecture and the scenario, and on the interaction of the decision makers executing complex tasks. These can be used to define surrogate measures for Speed of Command and Shared Situational Awareness; the accumulated delay time was used to compare the Speed of Command across architectures and the task window limit was used to evaluate Shared Situational Awareness. By using empirical data collected after a subject experiment, enhancements made to an existing model have resulted in a model that is more realistic and versatile in evaluating command and control architectures operating under different scenarios. The performance measures reflect the completion times of time critical tasks; including the variability of the decision maker's attention to the task, not only the initial delay in the Detect stage, but also delays that may occur due to interruptions in the task processing at the Select stage and the interaction of decision makers due to the synchronization of the launch of resources. The model will be a valuable tool for evaluating proposed future command and control centers.

References

[Boettcher and Levis, 1982] Kevin L. Boettcher and Alexander H. Levis, "Modeling the Interacting Decisionmaker with Bounded Rationality," IEEE Transactions on Systems, Man and Cybernetics, Vol. SMC 12, No. 3, pp. 334-344.

[Boettcher and Levis, 1983] Kevin L. Boettcher and Alexander H. Levis, "Decisionmaking Organizations with Acyclical Information Structures," IEEE Transactions on Systems, Man and Cybernetics, Vol. SMC 13, No. 3, pp. 384-391.

[Cothier and Levis, 1986] Phillippe H. Cothier and Alexander H. Levis, 'Timeliness and Measures of Effectiveness in Command and Control," IEEE Transactions on Systems, Man, and Cybernetics, Vol. SMC 16, No. 6, pp. 844-853.

[Diedrich et al., 2002] F.J. Diedrich, S.P Hocevar, E.E. Entin, S.G. Hutchins, W.G. Kemple, and D.L. Kleinman, "Adaptive Architectures for Command and Control: Toward An Empirical Evaluation of Organizational Congruence and Adaptation," Proceedings of the Command and Control Research and Technology Symposium, Monterey, CA.

[Handley et al., 1999] H. A. H. Handley, Z. R. Zaidi, and A. H. Levis, "The Use of Simulation Models in Model Driven Experiments," Systems Engineering, John Wiley & Sons, Inc, 2:108-128.

[Handley and Levis, 2003] H. A. H. Handley and A. H. Levis, "Organizational Architectures and Mission Requirements: A Model to Determine Congruence," Systems Engineering, John Wiley & Sons, Inc, Vol. 6, No. 3.

[Handley, 2003] Holly A. H. Handley, "Dynamic Modeling of the Strategic Studies Group Architectures," Memorandum, January 14, 2003.

[Hiniker, 2002] Paul J. Hiniker, "A Model of Command and Control Processes for JWARS." 7th International Command and Control Research and Technology Symposium, September 16-20, 2002, Quebec City, Canada.

[Levis, 1992] Alexander H. Levis, "A Colored Petri net Model of Intelligent Nodes," Robotics and Flexible Manufacturing Systems, J.C. Gentina and S.G. Tzafestas, Eds., Elsevier Science Publishers, The Netherlands.

[Levis, 1995] A.H. Levis, "Human Interaction with Decision Aids: A Mathematical Approach," in Human/Technology Interaction in Complex Systems, Vol.7, W.B. Rouse, Ed., JAI Press, 1995

[Levis et al., 1998] Alexander H. Levis, Holly Handley, Zainab R. Zaidi, "Five Stage Decision Maker Model," Memorandum, August 28, 1998.

[March, 1978] J.G. March, "Bounded Rationality, Ambiguity, and the Engineering of Choice," Bell J. Economics, Vol. 9, pp. 587-608, 1978.

[March and Simon, 1958] J.G. March and H. A. Simon, Organizations, John Wiley and Sons, NY.

[Remy and Levis, 1986] P.A. Remy and A. H. Levis, "On the Generation of Organizational Architectures using Petri nets," in Advances in Petri Nets, G. Rozenberg, Editor, Springer-Verlag, Berlin.

[Wohl, 1981] J.G. Wohl, "Force Management Decision Requirements for Air Force Tactical Command and Control," IEEE Transactions on Systems, Man, and Cybernetics, Vol. SMC-11, No.9.

Appendix A: Correlation Data and Statistics

For each trial in the subject experiment, indicated by Team, a correlation was performed between the experimental completion time data and the simulated model completion time data. This value, indicated by Correlation, was obtained by correlating the tasks that were completed by both the team and the simulation; this number varies by team and is indicated by Number of Tasks. For each correlation a significance test was performed by using the F statistic; the results of the test are shown in the columns F Value and F Significance. In all cases, the null hypothesis of no predictive value can be rejected.

Table A.1: Original Task Process Model Correlation Statistics Between Experimental and Simulated Output

Team	Correlation	Number of Tasks	F Value	F Significance
FSf2W	.77	62	89.60	1.64E-13
FSf2D	.83	59	124.14	6.16E-16
FSf2B	.75	60	76.02	3.87E-12
FMd2W	.81	58	107.95	1.12E-14
FMd2D	.88	48	165.78	7.37E-17
FMd2B	.92	57	315.21	1.95E-24
DSf2E	.90	52	217.20	7.88E-20
DSf2C	.84	54	127.74	1.28E-15
DSf2A	.88	57	195.66	9.26E-20
DMd2E	.84	61	141.15	2.74E-17
DMd2C	.95	62	602.50	5.50E-33
DMd2A	.97	60	822.47	5.96E-36

Table A.2: Enhanced Task Process Model Correlation Statistics Between Experimental and Simulated Output

Team	Correlation	Number of Tasks	F Value	F Significance
FSf2W	.71	70	70.76	3.91E-12
FSf2D	.75	74	90.29	2.44E-14
FSf2B	.68	74	60.53	3.94E-11
FMd2W	.79	75	119.56	4.94E-17
FMd2D	.87	63	183.15	4.97E-20
FMd2B	.92	72	405.64	7.67E-31
DSf2E	.82	63	124.71	2.21E-16
DSf2C	.76	66	87.10	1.49E-13
DSf2A	.67	70	55.39	2.25E-10
DMd2E	.79	71	114.08	2.88E-16
DMd2C	.88	78	253.89	6.14E-26
DMd2A	.90	77	334.47	2.30E-29