

Influence Net Modeling with Causal Strengths: An Evolutionary Approach

Julie A. Rosen, PhD[†] and Wayne L. Smith[‡]
Science Applications International Corporation
1710 Goodridge Drive, McLean, Virginia 22102

Abstract

An approach to investigating the human decision cycle, particularly that employed by individuals and organizations during crisis, is presented. The collaborative approach described here is especially beneficial in today's world of rapidly evolving, global situations within which U.S. security policies and operational plans are generated.

This paper continues the documentation of research in the field of *Influence Net* modeling. Specifically, we will address the capabilities required of an automated system to encourage and facilitate the collaboration, both real-time and evolutionary, of decision makers and their supporting experts. We present our research results that extend traditional Bayesian inference net structure to allow for interactive use by modelers unfamiliar with probability theory or who are unwilling to spend the excessive time required to specify the traditional Bayesian model. The results of this research, called Causal Strengths (CAST) Logic, have been implemented as software applications by the authors and their colleagues.

1 Introduction

Throughout history, humans have employed various methods for analyzing the behavior and rationale of opponents, whether they be adversaries in armed combat, competitors in financial endeavor,

or, or friendly supporters that require occasional persuasion to remain supportive. More recently, attempts have been made to extract academic and empirically-based expertise from decision makers and their supporting experts, and place it in an analytical framework. In many cases, this transformation has not succeeded because the gap between the decision making community and the framework modelers is too wide. This paper continues the documentation of research conducted by the authors in support of a collaborative decision making process [Rosen and Smith, 1994].

The remainder of this section presents the motivation for the authors' investigation. Section 2 provides a discussion of the terminology and definitions unique to the Causal Strength (CAST) logic. Section 3 presents the four-step process employed in the CAST algorithm. Finally, Section 4 summarizes the benefits of this evolutionary approach to inference net analysis.

1.1 Motivation for this Investigation

With the end of the bi-polar political world, decision makers in the U.S. national security arena are faced with an ever-increasing number of situations that have the potential to become crises. In this paper, the term crisis includes situations of economic instability, ideological or cultural contrasts, as well as the more traditional (and oftentimes military-based) political and diplomatic security concerns. These crisis situations may occur while the involved parties are at peace; however, crises left untended or inaccurately estimated tend

[†] Phone: (703) 556-7354; email: Julie.A.Rosen@cpmx.saic.com

[‡] Phone: (703) 448-6522; email: Wayne.L.Smith@cpmx.saic.com

towards armed conflict situations that affect U.S. national and global stability interests.

U.S. security decision makers, including military planners, no longer face a single national government opponent whose power derives primarily from its military's capabilities. Today, world "actors" capable of generating crisis and instability, perhaps unintended, also include individuals representing multi-national organizations and multi-national states. The behavior of this set of actors, and their attendant actions, expand the more traditional list of state-sponsored conflict situations. In addition, as technological advances make the "global economy" a reality, conflicts formerly considered "internal disputes" possess the ability to disrupt, even destroy, the processes governing everyday lives of the citizens of many nations. In recognition of these events, the U.S. security arena has expanded the military's roles and missions to include the following [DOD 1994]:

- Urban conflict—the insertion and extraction of forces, such as employed in Somalia;
- Distributed forces—insertion of forces, possibly deep insertion, such as currently deployed throughout the Bosnian theater; and
- Major regional conflict (MRC)—force-on-force deployment to a single theater or multiple, concurrent campaigns.

However, the characteristics of potential situations are not the only parameters that define U.S. security concerns. Budget realities that headline today's news also must be considered. As the 21st century approaches, the national security community increasingly is mandated to reduce the size of its infrastructure. Combat forces of the next decade will be significantly diminished—the size of U.S. forces as well as the numbers available from our traditional allies. Not only are the warfighting forces "taking the hit;" the planning and intelligence communities similarly are undergoing a reduction in force. In addition to the human factor, national security facilities are reduc-

ing their focus with attendant consolidation mandating the closure of bases—both CONUS and OCONUS. Similar financial constraints upon our allies are reducing the likelihood that "host-country basing" will be available when regional crises arise.

1.2 Statement of the Problem

As scenarios for crisis and conflict arise, members of the U.S. national security community are tasked to examine the behavior and capabilities possessed by both our allies and opponents. Traditionally, two categories of investigation and analysis have been employed to identify influence strategies and their operational implementations:

- Seminars, workshops, and informal communications that extract knowledge from experts in the field of study. Sometimes this information is captured in a paper report that presents the results of the study to the decision maker; typically, this capture is performed by a single member of the study group. However, whether or not the results of the knowledge elicitation are documented, the experts' underlying source material, assumptions, justifications, and reasoning are very rarely maintained. Such information is crucial not only for the current decision maker, but also for future decision makers and their analysts who require historical, empirical evidence as the situation evolves.
- Mathematical and computer-based models/simulations that attempt to estimate current and future states of "physics-based" phenomena. As with the first category, the results of this type of investigation usually are "watered down" for presentation to the consumer or other analysts. The input parameters and internal "rules" of such models are glossed over in the presentation of these results. Occasionally, these results are sufficient to estimate the status of a situation. However, as with the seminar technique, reducing the documentation and presentation of the model's underlying reasoning may lead to confu-

sion and misinterpretation by the decision maker. The problem is exacerbated when such models are revisited by future decision makers and their analysts.

Too often in the past, the traditional techniques for examining a situation have produced assessments that are not borne out in time. For example, workshop-like analysis indicates that a leader is not believed to possess aggressive intentions, but an unstable situation initiates because that leader is not in control of events.¹ Sometimes, both techniques are employed concurrently, producing conflicting results. For example, observable evidence and physics-based models/simulations prove that an adversary possesses the technical capability to conduct aggression, but the adversary “backs off” when an outside influence is applied.² In this case, the motivation, perception, and intentions of the adversary underscoring the resulting behavior may not have been accounted for correctly, if at all.

Moreover, in today’s world, technological advances that “speed up” the timeline towards crisis, and the proliferation of this technology to more and more actors, means that (potential) crisis situations will rapidly evolve. In addition, understanding these situations depends on a greater number of parameters that are not physics-based. The diversity of human motivation and perception must be addressed by today’s analysts responsible for identifying influence strategies and plans. Such diversity means that experts from an increasing number of domains must be included in the analysis process; for example, psychologists, historians, economists, international industrialists, diplomats, and philosophers.

However, as the community of analysts diversifies, problems of communication among these experts increases. Differences in terminology, knowledge, assumptions, and inference/reasoning

practices may lead to confusion and irreconcilable disagreement about the anticipated behavior and actions of a troublesome actor. Therefore, not only are a greater number of experts required to identify alternate influence strategies, but the decision maker must understand the interaction of parameters across domains of expertise in order to evaluate the alternate strategies.

1.3 *Influence Net* Modeling—A Collaborative Approach

In an attempt to address the goals of collaborative analysis, the authors have developed a technique for analyzing the causal relations of complex situations. This technique, known as *Influence Net* modeling, is a combination of two established methods of decision analysis: Bayesian inference net analysis originally employed by the mathematical community; and influence diagramming techniques originally employed by operations researchers. As illustrated in the following sections, *Influence Net* modeling incorporates both an intuitive, graphical method for model construction, and a foundation in Bayesian mathematics for the rigorous analysis of such models.

To facilitate communication among the domain experts and their “consumers,” *Influence Net* modeling encourages the domain experts, themselves, to create “influence nodes.” These influence nodes depict events that are part of (possibly complicated) cause-effect relations within the situation under investigation. The experts also create “influence links” between cause and effect that graphically illustrate the causal relation between the connected pair of events. This cause-effect relation can be either promoting or inhibiting, as identified by the link “terminator” (an arrowhead or a filled circle). The resulting graphical illustration is called the “*Influence Net* topology;” a sample topology is pictured in Figure 1.³

¹ For example, the fall of the Shah of Iran was not forecast sufficiently in advance, possibly due to the lack of analysis of the influence of non-political actors in that region.

² For example, the French Government applied pressure to Qaddafi just after the U.S. bombing of Libya. The power of French economic influence has been offered as one justification for Qaddafi’s response.

³ The sample *Influence Net* model was constructed from open source reference material investigating the influences on Saddam Hussein’s decision making after Iraq’s invasion of Kuwait in 1990 [Davis and Arquilla, 1991a, 1991b].

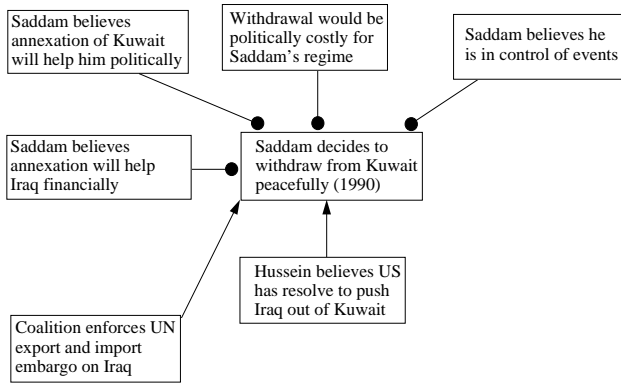


Figure 1. Sample *Influence Net* topology.

This topology of the *Influence Net* model is a significant product for presentation and communication of the current situation. In addition, the decision maker also requires a capability to examine courses of action that have the potential to prevent, mitigate, or contain crisis situations. In order to provide a real-time analytical capability, a rigorous mathematical foundation, such as Bayesian inference networks, was required.

In conducting this research, SAIC staff, with original ARPA sponsorship, developed an automated utility that assists intelligence analysts and operational planners in their examination of the influencing factors underlying the decision making cycle. The resulting Unix-based automated utility is called the Situational Influence Assessment Module (SIAM). A Windows-based version, called *Causeway*TM, is currently under development at SAIC.

During the functional requirements phase of the SIAM effort, SAIC staff consulted with (current and former) members of the intelligence community to identify their “working environment.” As a result, we discovered that although the traditional Bayesian inference network algorithms might be useful in assessing cumulative impacts of complex cause-effect relations, the majority of users of an automated system did not have the experience required to run the software models available. Moreover, even those users with a background in probability theory, or those modeling physics-based systems, often did not

have the resources required to define the complete transition matrix of conditional probabilities in order to begin their investigation. These two restrictions are exacerbated when the users of an automated application are policy makers involved with forecasting reactions in international affairs, which are continually changing.

In order to address these two issues, SAIC staff, in conjunction with members of George Mason University’s C3I Center for Excellence, developed an evolutionary approach to model construction—“Causal Strengths” (CAST). Quantitative parameters are specified by the user for the individual cause-effect relations that indicate the “strength” that the influencing event (cause) will promote or inhibit the modeled effect. In addition to being intuitive and accessible for the typical user, this approach is founded in the mathematics of robust Bayesian inference network. Therefore, although the typical user constructing an *Influence Net* model may not have the time, source material or experience required to provide the complete Bayesian model “from scratch,” the CAST Logic allows for the evolutionary incorporation of additional knowledge as it is obtained. After all, “[w]e can never trust what the old books tell us. We cannot even trust what our senses tell us... [o]nly reason can give us certain knowledge” [Gaarder, 1994]. The remainder of this paper describes this logic and its implementation within the SIAM automated application.

2. Causal Strengths Logic

As stated above, typical users of an automated collaborative utility will not have sufficient resources and/or experience to employ the traditional Bayesian inference model. Therefore, in developing the Causal Strength (CAST) logic, the authors and their colleagues [Chang *et al*, 1994] considered three goals:

1. Generate a user interface logic that requires a relatively small number of value assignments, which can be expanded to a full Bayesian model;
2. Permit the user to modify the model with parameters that are meaningful to them; and




Figure 2. Pairwise influence.

3. Produce a user-interface that follows a consistent inference logic understandable to the user.

The following paragraphs present the details of the CAST algorithm and its presentation to the user.

2.1 Causal Strengths Assignment

Heuristically, in saying that event **A** *causes* event **B** to occur, the speaker is implying that if **A** is present, it is likely that event **B** will result. English speakers using words such as *influence*, *cause*, *force*, and *mandate* are attempting to quantify the likelihood that **A** implies **B**. This quantification is called the *strength* of the implication.

In the simple pairwise cause-effect relationship illustrated in Figure 2, the strength of the implication is defined by a conditional probability. Note that in this simple example, the probability space “universe” is completely defined by these two events.

In probabilistic terms, the strength of causality, denoted h , can be viewed as the probability of an event **Z**, where **A** and **Z** fully determine **B**. For example, event **Z** is the result of a “flipped” coin. If the coin lands “heads up” ($h = 1$), then it is certain that the occurrence of event **A** will result in event **B**; if the coin lands “tails up” ($h = 0$), then **A**’s occurrence has no effect on event **B**.⁴ Again, this cause-effect relationship assumes that there are no other events in the universe; that is, “by itself,” event **A** implies/causes event **B**’s occurrence.

⁴ Note that a negative causal strength implies that the occurrence of event **A** will result in the *non-occurrence* of event **B**. This “inhibiting” causality is discussed in the sequel. The *magnitude* of the strength corresponds to the probability of event **Z** as described here.

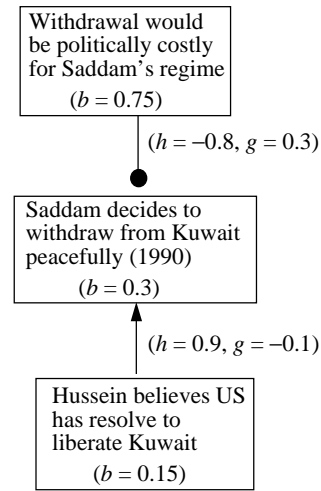


Figure 3. CAST parameters.

To complete the heuristic analogy, the “negation” of influencing events must also be incorporated. That is, for this simple illustration, the non-occurrence of event **A** implies (with a defined strength) that event **B** will occur anyway. Analogously, there is a second “coin flip,” say event **W**, whose likelihood, denoted g , determines the probability of event **B**; if $g = 1$, then the absence of event **A** is certain to produce event **B**’s occurrence; if $g = 0$, then event **A**’s non-occurrence has no effect on the outcome of **B**.⁵

In the preceding, we considered only pairwise causality. However, in modeling real world situations, domain experts need to include as many influencing events as time and experience allow. It is noted that since the entire “ground truth” universe is unknown, any modeling should be considered an approximation. The CAST logic is presented as a “first step” in the evolution of more refined approximations. It facilitates the generation of increasingly complete Bayesian inference models.

We illustrate these heuristic concepts with the *Influence Net* drawn in Figure 3.

Again, assuming pairwise causality, the truth in the event *Hussein believes US has resolve to*

⁵ Again, note that a negative strength implies that event **A**’s non-occurrence inhibits event **B**, thereby resulting in its non-occurrence.

liberate Kuwait causes the event *Saddam decides to withdraw from Kuwait peacefully* with a strength of $h = 0.9$. Conversely, if Hussein does **not** believe in the US resolve, then the causality has a lesser (inhibiting) strength of $g = -0.1$.

With this illustration we introduce two more concepts that complete the CAST logic construction: negative causal strengths, and baseline probabilities. The former is discussed here; the latter is introduced in Section 2.2. Again, heuristically, when we state that an influence *inhibits* the occurrence of an event, the implication is that the likelihood of the effect occurring is inversely affected by the presence of the influence. In the example of Figure 3, if the “parent node,” *Withdrawal would be politically costly for Saddam’s regime* were true, then Saddam’s decision to withdraw peacefully is **unlikely** to occur; specifically at a strength of $h = -0.8$. Note that if Hussein did **not** believe that withdrawal would be costly, then he still might consider withdrawal—although at a reduced strength level of $g = 0.3$. Taken together these two strengths imply that the overall influence of the political cost of withdrawal *inhibits* Saddam’s decision to withdraw from Kuwait. This inhibiting influence is designated with the filled circle terminator on the link.

2.2 Baseline Probabilities Assignment

Until now we have discussed only the pairwise causal relationship between two events. In order to account for the remainder of the “universe,” the notion of a *baseline probability* is introduced. An *Influence Net* model includes a finite set of cause-effect relationships; determined primarily by the time and resources available to the modeler. After identifying the cause-effect event pairs, the modeler assigns pairwise causal strengths to indicate the influence of the *parent* event “by itself” on the *child* event. As presented in the next section, the cumulative effect of the modeled influences is calculated for each affected event in the *Influence Net*. However, unidentified influences must also be included.

$X =$ *Saddam decides to withdraw from Kuwait peacefully.*

$Y =$ *Hussein believes US has resolve to liberate Kuwait.*

$Z =$ *Withdrawal would be politically costly for Saddam’s regime.*

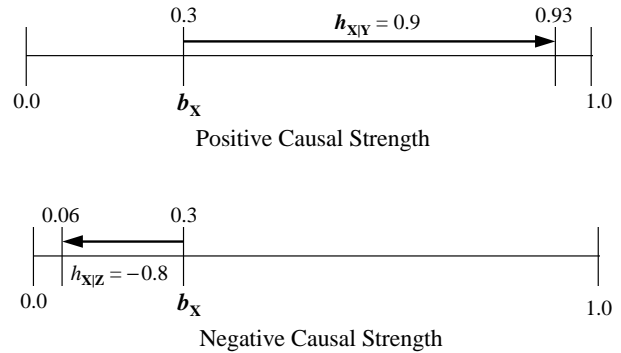


Figure 4. Causal strengths and baseline probabilities.

To account for this “remainder,” the baseline probability of an event (illustrated by nodes in our figures) is the user-assigned assessment that the event would occur **independent** of the modeled influences in the net. For the example depicted in Figure 3, all knowledge of the situation not explicitly included implies the likelihood of Saddam deciding to withdraw from Kuwait is $b = 0.3$.⁶

2.3 Combining CAST Parameters

To illustrate the method through which these causal strengths and baseline probabilities combine, consider the illustration in Figure 4. Here again, we use the sample model presented in Figure 3. For this graphical depiction, we assume that the two parent events, denoted Y and Z , respectively, are assumed to be true. Therefore, we consider their “true” causal strengths on the child event $X =$ *Saddam decides to withdraw from Kuwait peacefully*. In this figure the “true” causal strengths are denoted $h_{X|Y}$ and $h_{X|Z}$, respectively.

The event X has a baseline probability of 0.3; so, there is a 70% belief that Saddam will **not** decide to *withdraw from Kuwait*. That is, ignoring

⁶ Note that if we assume all information about the situation is included in the model, then there are no additional influences; i.e., the baseline probability of an event should be set to $b = 0.5$ in the absence of new information. As additional information relevant to this model becomes available, the baseline probability should be reconsidered and reassigned.

In the traditional full Bayesian inference network, event \mathbf{X} would be associated with a transition matrix of 2^n conditional probabilities. As stated throughout this paper, the CAST logic can be used to facilitate an “initial cut” at the daunting task of identifying the complete model. In many instances, this estimate is sufficient to support the decision making required, especially in today’s rapidly changing crisis situations. It is noted, however, that as modelers obtain more time and additional knowledge of the situation, the estimate can be refined until the more complete matrix is fully defined.

For the remainder of this section, we present the CAST algorithm for calculating the initial set of values for the transition matrix associated with each event. There are four steps that comprise the CAST algorithm. Once these steps are completed, the traditional probability calculations are performed to derive the cumulative likelihood of any event included in the *Influence Net*. The four steps that are performed for each conditioning case associated with the child event in the model are:

1. Aggregate positive causal strengths,
2. Aggregate negative causal strengths,
3. Combine the positive and negative causal strengths, and
4. Derive conditional probabilities.

To simplify the notation used in previous sections, let C_i denote the causal strength of the i^{th} parent event in the conditioning case of interest. For example, in Case 2 above:

$$\begin{aligned} C_1 = g_{\mathbf{X}|\mathbf{Y}} &= -0.1, & C_2 = h_{\mathbf{X}|\mathbf{Z}} &= -0.8, \\ C_3 = h_{\mathbf{X}|\mathbf{A}} &= -0.7, & C_4 = h_{\mathbf{X}|\mathbf{B}} &= -0.95, \\ C_5 = h_{\mathbf{X}|\mathbf{C}} &= +0.4, & C_6 = h_{\mathbf{X}|\mathbf{D}} &= -0.9. \end{aligned} \quad (3)$$

3.1 Aggregate Positive Causal Strengths

In this step, we combine the set of causal strengths with positive values for each of the 2^n conditioning cases associated with the selected child event. These positive causal strengths are aggregated using the combination rule in equation (4):

$$C_+ = 1 - \prod_i (1 - C_i) \quad \text{over all } C_i \geq 0 \quad (4)$$

This expression follows from the assumptions of independence of the pairwise cause-effect relationships. In other words, the complement of the aggregate positive causal influence C_+ (i.e., $1 - C_+$) is the increase in the probability that the child event will **not** occur unless (at least) one of the multiple independent influences cause the event to occur. For Case 2, above:

$$C_+ = 1 - (1 - 0.4) = 0.4 \quad (5)$$

3.2 Aggregate Negative Causal Strengths

Similarly, we combine the set of causal strengths with negative values for each of the 2^n conditioning cases associated with the selected child event. These negative causal strengths are aggregated using the combination rule in (6).

$$C_- = 1 - \prod_i (1 - |C_i|) \quad \text{over all } C_i < 0 \quad (6)$$

The interpretation for this rule is: The complement of the aggregate negative causal influence C_- (i.e., $1 - C_-$) is the increase in the probability that the child event will occur unless (at least) one of the multiple independent influences cause the event **not** to occur. For Case 2, above:

$$\begin{aligned} C_- &= 1 - (1 - 0.1)(1 - 0.8)(1 - 0.7)(1 - 0.95)(1 - 0.9) \\ &= 0.99973 \end{aligned} \quad (7)$$

3.3 Combine Positive and Negative Causal Strengths

In order to combine the aggregated positive and negative causal strengths, we introduce the following axiom:

Cancellation Axiom: Let $(1 - C_+)$ denote the potential of a child’s occurrence being promoted due to a set of parents, and let $(1 - C_-)$ denote the potential of a child’s occurrence being inhibited by a set of parents. Then, there is an overall influence, \mathbf{O} , that represents the net influence of the set of parents. The overall influence is given by the ratio of the aggregated promoting and inhibiting influences.

Heuristically, this axiom asserts that the accumulated influence of all parents (specified in the conditioning case) is partitioned into: a portion that balances out the “opposing side;” and the remaining overall influence.

From the form of this axiom, we now introduce the following expressions for calculating the overall influence on the child event, \mathbf{O} :

If $C_+ \geq C_-$, then solve for \mathbf{O} (≥ 0):

$$(1 - C_+) = (1 - \mathbf{O}) \times (1 - C_-) \quad (8)$$

If $C_+ < C_-$, then solve for \mathbf{O} (< 0):

$$(1 - C_-) = (1 - |\mathbf{O}|) \times (1 - C_+) \quad (9)$$

For Case 2 above, $C_- > C_+$, and:

$$(1 - 0.99973) = (1 - 0.4) \times (1 - |\mathbf{O}|) \quad (10)$$

which implies $\mathbf{O} = -4.5 \times 10^{-4}$.

Recall that this procedure is employed for each of the 2^n conditioning cases associated with the selected child event.

3.4 Derive Conditional Probabilities

The fourth, and final, step unique to the CAST algorithm follows from the cause-effect pair calculations introduced earlier. Specifically, the overall influence, \mathbf{O} , is used to update the baseline probability of the child event. The updated value then can be inserted into the traditional Bayesian transition matrix. (Again, recall that this procedure is performed for each of the 2^n elements of the transition matrix.)

Consider the j^{th} conditioning case. Let \mathbf{O}_j denote the overall influence on the child event from the j^{th} set of parent states. Then the conditional probability of the child, given the j^{th} set of parent states is given by:

$$P[\text{child} | j\text{th set of parent states}] = \begin{cases} b_{\text{child}} + (1 - b_{\text{child}}) \times \mathbf{O}_j; & \text{for } \mathbf{O}_j \geq 0 \\ b_{\text{child}} - b_{\text{child}} \times \mathbf{O}_j; & \text{for } \mathbf{O}_j < 0 \end{cases} \quad (11)$$

Again for Case 2 above, we have the following entry for the corresponding element of the transition matrix:

$$\begin{aligned} P[\mathbf{X} | \neg\mathbf{Y}, \mathbf{Z}, \mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{D}] \\ &= 0.3 - 0.3 \times (-4.5 \times 10^{-4}) \quad (12) \\ &= 0.300135 \end{aligned}$$

3.5 Calculating the Effect on the Child Event

Once the transition matrix is completed through the use of equation (12) on each of the 2^n parent states, then the traditional law of total probability can be employed to determine the current estimate for the likelihood of the child event. Specifically,

$$\begin{aligned} P[\mathbf{X}] = \\ \sum_{j=1, \dots, 2^n} P[\mathbf{X} | j\text{th set of parent states}] \times \\ P[j\text{th set of parent states}] \end{aligned} \quad (13)$$

Again, the typical analyst investigating a crisis will not have the resources to complete the joint probability matrix for real world situations. Therefore, a sufficient approximation is to assume the parents are independent. In Case 2 above, the joint probability, $P[j\text{th set of parent states}]$, on the right side of (13) is calculated as the product:

$$\begin{aligned} P[\neg\mathbf{Y}, \mathbf{Z}, \mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{D}] \\ &= P[\neg\mathbf{Y}] P[\mathbf{Z}] P[\mathbf{A}] P[\mathbf{B}] P[\mathbf{C}] P[\mathbf{D}] \quad (14) \end{aligned}$$

With the expert-assigned understanding of the situation as it existed in early 1990, the probability that *Saddam decides to withdraw from Kuwait peacefully* is less than 0.1. That is, the combined influences on Saddam Hussein do not serve to convince him to leave Kuwait peacefully.

4. Summary

Again we note that if the modelers had sufficient time and knowledge of the dependencies of the influences, then completion of the full transition matrix would produce the traditional results. However, as crisis situations evolve, uncertainty about the events' behavior does not permit this knowledge. Such uncertainty is exacerbated when the situation is not the “physics-based”

analytical problems that have been modeled in the past.

Of special importance in the evolution of such non-traditional models is the ability for groups of users in diverse fields of study to collaborate on portions of the model, bringing their complementary expertise to the whole network. Since the CAST logic requires relatively few parameters to initialize the model, automated implementations of this procedure can be employed in workshop/seminar settings to elicit knowledge about influencing relations. When such relations are documented in an automated system, such as **SIAM** and *Causeway*TM, sensitivity analysis becomes relatively easy to perform; thereby providing the workshop participants insight into the cumulative impact of the complex, and often conflicting, individual case-effect influences.

In this way, evolutionary refinement of the *Influence Net* model is possible and encourages consensus among the domain/knowledge experts. It is this goal—to capture knowledge directly from the domain experts—that is considered essential to the usefulness of an automated decision assistance utility in today's world of ever-changing influencing situations. The authors believe that the CAST logic, and its **SIAM** and *Causeway*TM implementations serve as a successful first step in the evolutionary process known as distributed, collaborative decision making.

References

[Chang *et al*, 1994] K.C. Chang, Paul E. Lehner, Alexander H. Levis, S. Abbas K. Zaidi, and Xinhai Zhao. *On Causal Influence Logic*, Technical Report for Subcontract no. 26-940079-80, George Mason University, Center of Excellence for C3I, December 3, 1994.

[Davis and Arquilla, 1991a] Paul K. Davis and John Arquilla. *Deterring or coercing opponents in crisis: Lessons from the war with Saddam Hussein*. Rand Corporation, Santa Monica, California, 1991.

[Davis and Arquilla, 1991b] Paul K. Davis and John Arquilla. *Thinking about opponent behavior*

in crisis and conflict: A generic model for analysis and group discussion. Rand Note N-3322-JS, Rand Corporation, Santa Monica, California, 1991.

[DOD 1994] Director, Defense Research and Engineering. *Defense Science and Technology Strategy*. Department of Defense, Washington, D.C., September 1994.

[Gaarder, 1994] Jostein Gaarder. *Sophie's World*. Berkley Publishing Group, New York, NY, 1994. Translated by Paulette Møller.

[Rosen and Smith, 1994] Julie A. Rosen and Wayne L. Smith, Cooperative Security Issues: Assessing the Effect of Situational Influence Measures. In *Proceedings of 1994 Symposium on C2 Research and Decision Aids*, pages 510–524, Monterey, California, June 1994. Joint Directors of Laboratories.