RC24511 (W0803-009) March 3, 2008
Other

IBM Research Report

Using Stochastic Process Algebra Models to Estimate the Quality of Information in Military Sensor Networks

David J. Thornley[1], Chatschik Bisdikian[2], Duncan F. Gillies[1]

[1]Department of Computing
Imperial College London
180 Queen's Gate
South Kensington, London
UK

[2]IBM Research Division
Thomas J. Watson Research Center
P.O. Box 704
Yorktown Heights, NY 10598
USA

Using stochastic process algebra models to estimate the quality of information in military sensor networks

David J. Thornley[a], Chatschik Bisdikian[b] and Duncan F. Gillies[a]

[a]Department of Computing, Imperial College London,
180 Queen's Gate, South Kensington, London, UK
djt@doc.ic.ac.uk,dfg@doc.ic.ac.uk;

[b]IBM Corporation, Thomas J. Watson Research Center
19 Skyline Drive, Hawthorne, NY, USA
bisdik@us.ibm.com

Keywords: quality of information, sensor networks, timeliness, stochastic modeling, network performance

ABSTRACT

In a typical military application, a wireless sensor network will operate in difficult and dynamic conditions. Communication will be affected by local conditions, platform characteristics and power consumption constraints, and sensors may be lost during an engagement. It is clearly of great importance to decision makers to know what quality of information they can expect from a network in battlefield situations. We propose the development of a supporting technology founded in formal modeling, using stochastic process algebras for the development of quality of information measures. A simple example illustrates the central themes of outcome probability distribution prediction, and time-dependency analysis.

1. INTRODUCTION

In military applications, command demands information with defined levels of accuracy and confidence, gathered through sensor deployments and cued actions commonly specified beforehand. This may relate to long-term strategic planning and activity (*e.g.* a humanitarian operation), shorter term tactical field operations (a day's reconnaissance), or to short duration, urgent events such as the tracking of incoming rocket-propelled weaponry, or detection of a sniper's gunshot. In all these situations, the effectiveness of the decision making depends on the *quality of information*, or "QoI" available.

Traditionally, QoI has been studied in relation to the collection of information, its storage it in data warehouses, and the efficacy of its retrieval. Important quality attributes include accuracy, consistency, relevancy and timeliness.[1–3] In a military context, the decision making opportunities created by just-in-time, sensor-data-rich information environments are opening up new and challenging research directions. QoI may be used to characterize the standard of the data flowing through the sensor network, and of the information derived from processing these data.[4] We anticipate that the use of QoI attributes in these environments will provide an efficient, effective means for assessing the applicability of sensor-derived information to missions.

It would be exciting to find general attributes which could be assigned to given pieces of equipment that result in effective fusion under all circumstances. However, the same piece of information in different contexts may or may not be of value. For example, a low-quality, grainy, infrared photograph revealing persons moving along a monitored path could be sufficient for the detection of important activity, but not for identifying and

Research was sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence and was accomplished under Agreement Number W911NF-06-3-0001. The views and conclusions contained in this document are those of the author(s) and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

classifying the materiel under the group's control. This examples indicates that quality of information is specific to the combination of devices and mission goals.

For the purpose of this work, we consider that sensor-derived information is used to create a report of the state of an entity in the scenario. Such a report may be inaccurate, and so we are interested in the difference between the reported and actual state. For example, if the mission is to locate and monitor vehicles or targets, measurements of acoustic, seismic, visual or radar features may be gathered and fused to aid in detecting the presence of targets, their number, location (to some specified accuracy), and approximate velocity vectors. Bisdikian[4] discusses detection and tracking through a multi-modal network of sensors, with measurement tolerances that can influence QoI, and suggests that suitably specified QoI metrics could drive the autonomic operation of sensor networks. These would adjust network-derived QoI values in the presence of changing mission needs or sensor network capabilities.

Field observations can be affected by limitations of the sensing channel (e.g., foliage obscuring or falsely cueing optical sensors), the sensing operation itself, and the communications path between the sensor and the fusion center. As a result, field observations will provide an imperfect representation of the situation, giving rise to uncertainty in estimated parameters. Such uncertainty is commonly represented by a probability distribution or density function (pdf) over parameter values. These pdf's also exhibit time-dependent behavior, capturing the evolution of uncertainty in the parameters until, for example, a new set of field observations is fused into our current impression of the facts.

It is an objective of our research endeavor to introduce modeling and analysis techniques for sensor-enabled missions that will allow us to quantify the time-evolution of uncertainty in important parameters. This in turn will permit us to estimate the QoI expected from variant configurations of the sensor deployments. Modeling of this kind will therefore be of significant benefit to selecting the design, and optimization of the settings, of military sensor networks embedded in the context of different operational plans and scenarios. We hope eventually to directly synthesize design regimes, quality of information measures and fusion schemes. In this initial work, we provide an introductory example demonstrating that information quality measures can be calculated from solutions to probabilistic models, and illustrate the advantages of doing so using process algebra.

2. STOCHASTIC MODELING AND PROCESS ALGEBRAS

When a system involves many components with rich *inhomogeneous* structure and interactions, it is helpful – and generally necessary – to define components whose individual behaviour we can comprehend, then compose them together to generate the interactions which lead to the over-all behaviour of the system. When we need access to probability distributions over the state of the system, this is approached in using a process algebra. Process algebras are defined, analysed and implemented with the goal of producing formally verifiable, transformable and comparable models. Timed stochastic process algebras (SPAs) are generally intended for use in modeling real systems with temporal behaviours, such as the present monitoring example. We therefore explore the use of distribution-based models defined using SPAs, because this keys into a substantial literature providing formalisms (*e.g.*[5,6]) and tools for solution and model checking (*e.g.*[7–9]) which will assist in development, validation and generalization.

A wide range of SPAs have been proposed and explored in the literature, each intended to address a particular aspect of behaviour well. The "Performance Evaluation Process Algebra", or PEPA, introduced by Jane Hillston[5] was designed to address the needs of the performance analysis community to analyse interacting – or cooperating — processes in a continuous time Markovian framework. Other process algebras are motivated by other aspects of a general system. For example, Ambients[10] are designed for the purpose of modeling activities in hierarchical control volumes. The stochastic Pi calculus on the other hand, models communications between processes explicitly. We have initiated the present work in PEPA because of our interest in the interactions between sensor deployments and environmental conditions in the form of cooperations. As the work progresses we will take advantage of the formally provable translations between popular calculi, and their solution in maturing tool sets, to optimize the degrees of freedom in our models.

We model a sensor network in a process algebra by specifying a component process for each sensor, network node, environmental aspect and actor. These comprise behavioural states, and incorporate actions with associated execution time distributions leading to transitions between those states. In our example, these mimic a sensor package (combination of sensor and network node selection and placement) in a deployment scenario including environmental conditions which affect the functioning of the infrastructure. The state of the whole system (comprising the state currently occupied by each component) then moves around a chain of states according to the actions and combination of actions. This is most commonly a Markov chain, for which a rich set of mathematical manipulation and solution tools exist. This Markov chain is represented as a vector of states and a transition matrix of Poisson rates.

Each action follows a probabilistic description of its duration in the model. The resulting stochastic model captures the *potential evolutions* of the system as a locus annotated with a probability measure. For example, the model could provide a probability distribution over the potential locations of a moving target. This enables the specification of a time-dependent evolution of entities interacting over specified actions. For example, a sensor may not detect an acoustic event if it is masked by its acoustic environment. In the model we show later, we create this behaviour by requiring the sensor, the target, and the acoustic environment to *cooperate*over the hearing action. If one of these three components does not make the hearing action available, it will not occur.

With a suitably constructed model, we may examine selected outcomes, or groups of outcomes, in terms of their probability during some form of equilibrium or mean behaviour, or as a transient response. An example equilibrium behaviour may represent the distribution of target locations in an on-going monitoring scenario, and perhaps specifically coinciding with certain events. For example, the location distribution at the arrival of a tracking update from a sensor network will provide a measure of the predictable accuracy of such an update. A transient response takes a starting state description, and provides the state distribution at a set instant in time thereafter. For example, we may calculate the probability distribution of a target location as a function of time in the period between updates from a sensor network.

2.1 PEPA

We are taking the initial steps in this work in PEPA[5] due partly to its compatibility with our desire to analyse cooperating processes, and partly to the availability of directly applicable solution tools, and translators from PEPA into input forms for tools which analyse aspects of a system beyond the defined scope of PEPA itself, such as transient behaviours.

A system model is constructed in PEPA by defining one or more component processes, each having one or more states between which that process may transition as Poisson processes of constant mean rate. Any transition in a process may be defined to cooperate – *i.e.* be constrained to occur simultaneously with – a transition in one or more processes with the same label. If two objects cooperate over an action labelled "action", but one defines the rate as a, and the other defines the rate as b, then the cooperation will proceed as a Poisson process of rate $\min(a, b)$. The cooperation formalism can be used to throttle or stop an activity by introducing a controlling component. In the example we detail below, the acoustic environment can limit the ability of the sensor to detect the target.

At least one of the actions in a cooperation set must have a finite Poisson rate. In a cooperation with unequal rates defined for the actions involved, the required joint transition occurs as a Poisson process with the minimum rate taken from those defined in the cooperation.

3. STOCHASTIC MODELING OF QUALITY OF INFORMATION

Our overall goal is to produce a function which takes a description of a scenario, and provides one or more QoI measures. Here, we construct a small, simplified example to produce a measure of information accuracy, and an illustration of the decay of the quality of that information over time.

The quality of information embodied in a message may be interpreted as a probabilistic constraint on the outcomes we should expect based on having received that message. If we receive a message with the essential content "the target is in range", then the *prima facie* interpretation is that the target is in range, with probability 1. In military situations in particular, there is no such thing as certainty, so we want to deal with quantitative measures of the relationship between an apparently decisive declaration embodied in such a signal, and the outcomes we expect in the field when it is received.

A scenario is characterized in a model, which may be empirical (statistical results of field tests), intuitive (logical constraints on behaviours based on experience with equipment and personnel), speculative (creating modeling factors which, if feasible, will be beneficial) or formal (abstract models in a suitable algebra). The work in this paper forms the first step in the formal modeling aspect which is to be integrated in a holistic manner with the whole range of perspectives, creating a framework within which the necessary mathematical links between practice, theory and experience will be formulated.

The actions in a Markov chain model proceed as Poisson processes, which have a fixed form of probability distribution function (pdf) over the *time between transitions* or *sojourn time in the source state* with one parameter: the rate constant. The pdf of time delays between Poisson events is $te^{-\lambda t}$, where t is the free variable time, and λ the parameter we may set. This distribution decreases monotonically. Models based on continuous time Markov chains (CTMCs) are commonly employed in the estimation of mean behaviours. However: few real processes follow the Poisson distribution exactly, so higher moments of delays are poorly reproduced. This is remedied by substituting a clique of states for a single state, which enables approximation of more realistic time distributions by a suitable choice of transition rates within that clique (*e.g.*[11]). However, for the measure we illustrate here, mean behaviour is an acceptable approximant, and leaves the model structure simple.

4. A SIMPLE MONITORING SCENARIO

Consider an actor who must be made aware of the location of a target which can be either "here" represented as occupying location 1, or "away" represented as occupying location 2. We place a sensor and associated network node in each location, and allow the network nodes to communicate with each other to transport detection reports from each sensor. The actor receives reports from her local network node which may have originated from either location's sensor. The target moves between the two locations independently. The sensors detect the presence of the target acoustically. To begin introducing realism, we include an acoustic environment which can block detection of the target if ambient noise levels are high enough. The actor, sensors, network nodes, target and acoustic environments are each modeled as individual components, which are combined through PEPA cooperations, indicated in Figure 1.

Each state in the Markov chain derived from Figure 1 is a *joint* state of all the component processes. For example, one such joint state will represent that both sensors are active, but not sampling, the target in location 1, the actor believing the target to be in location 2, the network node in location 1 ready to send to the node in location 2 that the target is in location 2, and both acoustic environments loud. This one of the states less fortunate for the actor, and indeed we can calculate the probability of being in that state.

When the actor receives a report specifying the location of a target which is not under their control, she updates her opinion of where that target is. The most trivial update would be to adopt the *prima facie* interpretation of the report: the target is now where this report states it to be. However, sensing, data transmission and fusion are imperfect, so the behaviours leading to that imperfection must be modeled to enable prediction of an interpretation with quantifiable attributes pertinent to the intended use of the information; *i.e.* the quality of information.

When a network node receives a report from its local sensor or from another node in the network, it immediately prepares to send that information to the larger network. This immediate retransmission of the information to the wider network is chosen to enable simple extension of the number of locations in our model in further work by creating a ring or star connectivity. A sensor's detection of the target will tend to be reported to all locations in such a system. This is design from the point of view of expressivity of the model, which is to be augmented with features drawn from practical network protocols.

4.1 Monitoring scenario PEPA components

We define the monitoring scenario shown in figure 1 as a set of components to be composed together through cooperations over certain actions. The state of the system as a whole comprises the product of the component state spaces. When we refer to a state which comprises $A1at1$, it can be any of those joint states which includes $A1at1$ in its description.

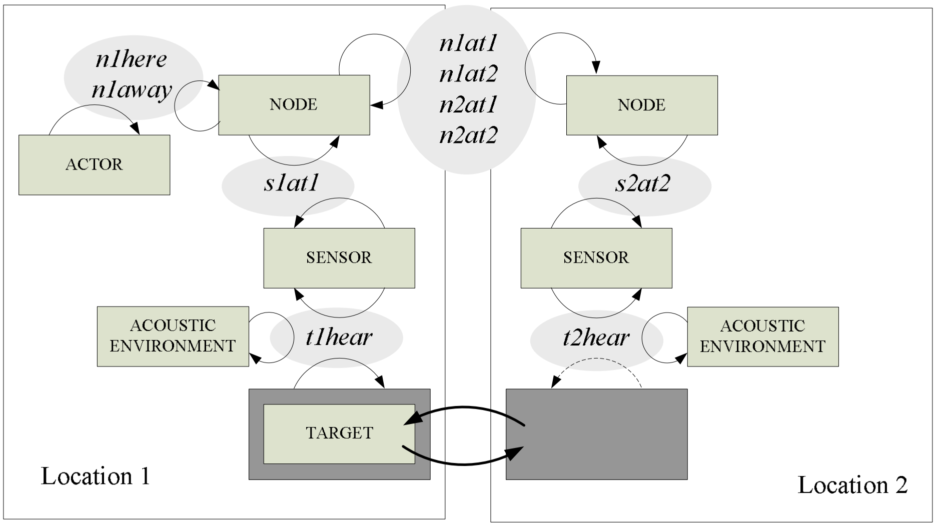

Figure 1. The components of our system follow a pattern of state transitions, some of which result from actions *cooperating* with actions in components. Components are shown boxed, and actions over which components cooperate are shown as shaded ellipses. For example, the sensor in location 1 has an action $t1hear$ which cooperates with the identically labelled action available in the target model, but only when it is in location 1, (*i.e.* susceptible to detection by the sensor), and the acoustic environment does not mask it (where "not masking" is modeled by making the transition available in the cooperation). The sensor in location 2 will not detect the target, because it is not present to satisfy the cooperation on the action. $t2hear$.

4.1.1 The actor

Our actor is stationed in location 1 and receives information from the network about the target through the $n1here$ or $n1away$ actions which assert respectively the presence or absence of the target. State $A1at1$ indicates that the actor believes the target to be present in location 1, and $A1at2$ believes it to be away in location 2.

$$\begin{aligned} A1at1 &\stackrel{def}{=} (n1here, n1hererate).A1at1 \\ &+ (n1away, n1awayrate).A1at2; \\[1em] A1at2 &\stackrel{def}{=} (n1here, n1hererate).A1at1 \\ &+ (n1away, n1awayrate).A1at2; \end{aligned}$$

4.1.2 The target

The target moves between locations 1 and 2 independently, and may be heard by a sensor in the same location.

$$\begin{aligned} T1Active &\stackrel{def}{=} (t1hear, t1hearrate).T1Active \\ &+ (t1activemove12, t1activemove12rate).T2Active; \\[1em] T2Active &\stackrel{def}{=} (t2hear, t2hearrate).T2Active \\ &+ (t2activemove21, t2activemove21rate).T1Active; \end{aligned}$$

State $T\{x\}Active$ indicates the target is in location x. The activity is specified to allow for inclusion of states in which it may, for example, be silent or in a stealth mode. Actions $t\{x\}activemove\{xy\}$ give the rate of the process of moving from location x to location y. The action $t\{x\}hear$ describes, from the point of view of the target, the process of *being heard*.

4.1.3 Sensors

The sensors operate independently, one located in each of locations 1 and 2, and hear the target only if it is in the same location and it is sampling. The transition $t1miss$ is included to avoid deadlock in this simple model, but may also be used as a handle on false negatives. The sensor in location 1 is defined as follows:

$$
\begin{aligned}
S1Active &\stackrel{def}{=} (s1sample, s1samplerate).S1Sampling;\\[1em]
S1Sampling &\stackrel{def}{=} (t1hear, t1hearrate).S1Heard\\
&\quad + (t1miss, t1missrate).S1Active;\\[1em]
S1Heard &\stackrel{def}{=} (s1at1, s1at1rate).S1Active;
\end{aligned}
$$

The sensor in location 2 is defined similarly, making textual replacements in the action and rate labels of 1 with 2 and *vice versa*.

$$
\begin{aligned}
S2Active &\stackrel{def}{=} (s2sample, s2samplerate).S2Sampling;\\
S2Sampling &\stackrel{def}{=} (t2hear, t2hearrate).S2Heard\\
&\quad + (t2miss, t2missrate).S2Active;\\
S2Heard &\stackrel{def}{=} (s2at2, s2at2rate).S2Active;
\end{aligned}
$$

4.1.4 Network nodes

The network comprises a node in each location which registers a detection event if the sensor makes a $s1at1$ transition, and prepares to send a message. In this network of two nodes, the structure is symmetric, with the node in location 1 defined as follows, then the node in location 2 defined by replacing 1 with 2 and vice versa:

$$
\begin{aligned}
N1Receive1 &\stackrel{def}{=} (s1at1, s1at1rate).N1Send1\\
&\quad + (n2at2, n2at2rate).N1Send2\\
&\quad + (n1here, n1hererate).N1Receive1;\\[1em]
N1Receive2 &\stackrel{def}{=} (s1at1, s1at1rate).N1Send1\\
&\quad + (n2at1, n2at1rate).N1Send1\\
&\quad + (n1away, n1awayrate).N1Receive2;\\[1em]
N1Send1 &\stackrel{def}{=} (n1at1, n1at1rate).N1Receive1\\
&\quad + (s1at1, s1at1rate).N1Send1\\
&\quad + (n2at1, n2at1rate).N1Send1\\
&\quad + (n2at2, n2at2rate).N1Send2\\
&\quad + (n1here, n1hererate).N1Send1;\\[1em]
N1Send2 &\stackrel{def}{=} (n1at2, n1at2rate).N1Receive2\\
&\quad + (s1at1, s1at1rate).N1Send1\\
&\quad + (n2at1, n2at1rate).N1Send1\\
&\quad + (n2at2, n2at2rate).N1Send2\\
&\quad + (n1away, n1awayrate).N1Send2;
\end{aligned}
$$

The network node maintains a representation of which location the target is in (the number appended to the state name) which is decoupled from that of the actor. For example, we could be in a state which comprises $A1at1$ and $N1Receive2$ (we refer to a system state which includes a particular sub-state as part of its description as *comprising* that sub-state). This allows in a simple manner for the actor to be misinformed even if the network's record is correct. This network node is constructed such that we could easily grow the network to a star configuration, or a ring, despite the tiny state structure.

Action $s1at1$ cooperates with the sensor in the same location to indicate that the target has been sensed locally. Action $n2at2$ cooperates with the node in location 2 to pass a message from that node indicating that it believes the target to be in location 2. The node in 2 may also indicate that it believes the target to be in 1 - this is an arbitrary choice to leave design options open.

When the node in location 1 is ready to send a message to the node in location 2, it enters state $N1Send1$ if it is to indicate the target is in location 1, and $N1Send2$ if in location 2. The belief state may be updated by the local sensor at any time by a $s1at1$ action, or to either belief state by a message from the other network node. This behaviour is simplified, and intended to allow for simple extension in subsequent modeling experiments.

Actions $n1here$ and $n1away$ inform the actor (through cooperation) that the target is "here" (in location 1) or away (in location 2). We could code these actions as passive in either the node or the actor, indicating respectively the actor polling the node, or the node emitting updates automatically.

4.1.5 Acoustic environments

We introduce an acoustic environment to intermittently block the sensors from hearing the target, independently of their sampling behaviour. It is either loud or quiet; either blocking or allowing detection by the sensor. For location 1, this is modeled as follows, using $AE0in1$ to indicate zero noise in location 1, *i.e.* quiet, and $AE1in1$ to indicate some noise in location 1, *i.e.* noisy. An environment transitions between levels of noise according to independent actions $acoustic\{xy\}ratein\{z\}$ jumping from level x to level y at the specified rate in location z.

$$
\begin{aligned}
AE0in1 &\stackrel{def}{=} (acoustic01in1, acoustic01ratein1).AE1in1 \\
&\quad + (t1hear, t1hearrate).AE0in1; \\[1ex]
AE1in1 &\stackrel{def}{=} (acoustic10in1, acoustic10ratein1).AE0in1;
\end{aligned}
$$

The independently defined process for the acoustic environment in location 2, we replace "in1" with "in2". When we include these environments in the model cooperating over actions $t1hear$ and $t2hear$ with the target, a sensor only hears the target if it is present, and the acoustic environment enables the corresponding hear action. This model can be extended to include more levels of noise, and enriched to cause gradual reduction of the ability to here the target by associating low rates with the hearing action to reduce the likelihood of it completing before the sensor registers a miss.

4.1.6 The monitoring scenario system

The whole system is formed by composing the component processes together through cooperations. In PEPA, this is made explicit through cooperations over specified actions, which is notated as a bow-tie ($\underset{\{action1, action2, \ldots\}}{\bowtie}$), and the set of actions written under the bow-tie. If no set of cooperation actions is specified, the processes are independent, running in parallel.

The state of each component of the system given in the definition is chosen so as to provide a notional starting point for evolution of the system. This ensemble must represent a feasible state for the total system, *i.e.* one which will be revisited, and can be reached by all other operating states. The whole system is shown here with some bracketed cooperating pairs laid out vertically to simplify presentation:

$$
System \stackrel{def}{=} A1at2 \underset{\begin{Bmatrix} n1here, \\ n1away \end{Bmatrix}}{\bowtie} \left(\left(\begin{array}{c} N1Send2 \\ \bowtie \\ \begin{Bmatrix} n1at1, \\ n1at2, \\ n2at1, \\ n2at2 \end{Bmatrix} \\ N2Receive2 \end{array} \right) \underset{\begin{Bmatrix} s1at1, \\ s2at2 \end{Bmatrix}}{\bowtie} \left(\left(\begin{array}{c} S1Active \\ \underset{\emptyset}{\bowtie} \\ S2Active \end{array} \right) \underset{\begin{Bmatrix} t1hear, \\ t2hear \end{Bmatrix}}{\bowtie} \left(\left(\begin{array}{c} AE1in1 \\ \underset{\emptyset}{\bowtie} \\ AE1in2 \end{array} \right) \underset{\begin{Bmatrix} t1hear, \\ t2hear \end{Bmatrix}}{\bowtie} T1Active\right)\right)\right)
$$

Reading the system equation from the right, we see the target cooperating with the acoustic environment, then these cooperating with a sensor to represent the sensor hearing the target if the acoustic environment allows for it. The sensors pass on reports to their own network node because the actions are named specifically for the locations. The network nodes share their information with each other, and the node in location 1 passes it on to the actor.

5. EQUILIBRIUM AND TRANSIENT SOLUTIONS FOR UPDATE ACCURACY

We want to calculate the probability distribution over potential locations of the target at the instant of arrival of a message from the local network node stating that it is present, specifically changing the actor's opinion from "the target is at location 2" to "the target is at location 1". In the PEPA model, this corresponds to an $n1here$ transition from any state of the whole system which has the actor in state $A1at2$ *i.e.* believing the target to be in location 2, from *any* state which comprises $A1at1$.

Solution for the equilibrium state occupation probabilities of the system is a fundamental result when analysing a Markov chain. The majority of interesting results are constructed using its probabilities, includng the rates of the reversed process,[12] as will be of use in the example calculation below. To construct the probability distribution we want here, we consider the circumstances in terms of the $n1here$ process. This is a Poisson process, so on entry to a state which is a target of that transition, we can take the subset of states which are reached by that action and normalise their probabilities from the total equilibrium (we can do this essentially because of the random observer property[13]).

5.1 Calculation procedure

We consider an observation taken at an independent instant, seeing state S_i. The state occupation probabilities in the continuous Markov chain of the model at an independent instant (*i.e.* one chosen without reference to system state) conform to the equilibrium distribution found by solution of the Chapman Kolmogorov probability flux balance equations.[13] Let the equilibrium state occupation probability of state S_i thus calculated be e_i.

Define F as the set of indices in list of states S which comprise $A1at1$. Define H similarly for states comprising $T1Active$, *i.e.* the target being in location 1 in any state of operation. For states we are interested in, the observed state is indexed in F, and the *previous* state *not* indexed in F. We then want the probability p_1 of such observed states being indexed in H.

If we can calculate a value $\hat{p}_1$ for the probability of the target being in location 1, we can compare the distribution of locations of the target ($p_1 = \hat{p}_1$, $p_2 = 1 - \hat{p}_1$) at the report arrival instant with that implied by the signal ($p_1 = 1$, $p_2 = 0$). This probability $\hat{p}_1$ is to be conditioned on the previous state having comprised $A1at2$. This is the probability of the observed current state having been entered from $S_j \notin F$.

To calculate this, for each state $S_i, i \in F$, we sum the probability fluxes of the *reversed process*[12] $f_{j,i}$ between that state and states $S_j, j \notin F$ to give f_i, and the probability fluxes in the reverse process $g_{j,i}$ between it and states $S_k \in F$ to give g_i. The probability c_i of observing the state S_i when it was previously a state $S_j, j \notin F$ is then $c_i = \frac{f_i}{f_i+g_i}$. This is the proportion of the reverse process' probability flux into a state comprising $A1at2$ from a particular observed state comprising $A1at1$. Multiplying the forward process' rate from state i to state j by e_j and dividing by e_i gives the rate in the reverse process from state j to state i. The reverse process has the same equilibrium solution as the forward process.

The modeled probability $\hat{p}_1$ of the target being in location 1 at the instant of arrival of a message represented as $n1here$ from a state $S_j, j \notin F$ is then:

$$\hat{p}_1 = \frac{n_h}{n_h + n_a}, \text{ where } n_h = \sum_{i \in H} e_i c_i,\ n_a = \sum_{i \notin H} e_i c_i, \text{ and } c_i = \frac{f_i}{f_i + g_i}$$

This is the probability that a signal indicating that the target is in location 1, given to the actor in location 1 who currently believes the target to be in location 2, is correct.

The results in table 1 are for all rates in the model set to 1.0, except the target motion rate of 0.01, the hear action rate of 100.0, and the rate at which the acoustic environments return to silence, which is set to 10.0, except where varied in the results table.

varied environment(s)	Varied rate: Update	0.1	1.0	10.0	100.0
1 and 2	2 to 1	0.8966	0.9654	0.9734	0.9822
	1 to 2	0.9036	0.9728	0.9814	0.9896
2	2 to 1	0.5976	0.9522	0.9734	0.9806
	1 to 2	0.9962	0.9881	0.9814	0.9898
1	2 to 1	0.9934	0.9819	0.9734	0.9823
	1 to 2	0.5376	0.9614	0.9814	0.9880

Table 1. $\hat{p}_1$ for update 2 to 1 ($A1at2$ to $A1at1$) and $\hat{p}_2$ for an update 1 to 2 for a range of values of $acoustic10ratein1$ and $acoustic10ratein0$. A lower rate increases the proportion of time spent in a state which does not permit the sensor in the same location to hear a target. The three groups of results are for the environments in both locations varied simultaneously, then constant at 10.0 in location 1 and varied in 2, then constant in 2 and varied in 1.

The rates used to generate these results are contrived to show that different parts of the system can interact in surprising ways. We would expect the probability of the target being where the network reports it to drop with higher noise levels, and this is the case when we vary the acoustic environments symmetrically. However, we can pick rates for the system which lead to counter-intuitive results, as seen in the non-monotonic response when the environments are differentially varied. This means that an intuitive result from a single test run of a system should not be regarded as proof that the behaviour will follow expectations under all conditions. Part of our research path will include automating the discovery of counter-intuitive behaviours, whether beneficial or detrimental to a mission. These features then motivate re-design of the model to improve generality.

Augmentation of the action behaviours will enable discovery of any counter-intuitive effects like this in real systems.

5.2 Time evolution of meaning

To calculate the evolution of the target location probability distribution over time after an update, we look again to the probability flux. This flux describes the rate of change of probability of the states as a redistribution or flow of probability between them. This gives us ordinary differential equations, which we integrate with respect to time, starting from the probability distribution at the update instant.

Specifically, we are interested in the probability distribution over potential locations of the target which ought to be assumed if no further information arrives. The arrival of this information does not affect the target's motion, because there is no feedback from the actor to the target, so in this specific example model, we can directly calculate the target's location independently (we are effectively selecting those state evolution sequences which do not include arrival of a signal). We show the time evolution of the meaning of an update $A1at2$ to $A1at1$ via action $n1here$ in terms of the change in p_1, denoting $t1activemove12rate$ and $t2activemove21rate$ as α, since they have the same value in our example:

$$p_2(t) = 1 - p_1(t)$$

$$\frac{dp_1}{dt} = \alpha(p_2 - p_1)$$

$$\frac{dp_2}{dt} = \alpha(p_1 - p_2)$$

Integrate simultaneous ODEs from $p_1(0)$ and $p_2(0)$ over time t

$$\Rightarrow p_1(t) = \frac{1}{2}\left(p_1(0)(1 + e^{-2\alpha t}) + p_2(0)(1 - e^{-2\alpha t})\right)$$

This begins with the distribution we calculated at update, and decays to the equilibrium probability. This expression is only valid for this particular model, but we will see this general decay trend in a number of characteristics of interest. These could be referred to as information quality decay or staleness, and knowledge of the precise geometry of such curves will be essential in properly calculating their effect in a scenario.

6. CONCLUSIONS

We have demonstrated that modeling a target monitoring scenario using a stochastic process algebra allows us to quantify aspects of the quality of disseminated information depending on the environment, target, sensor and network behaviours in a time-dependent manner. We are by no means unique in desiring an effective treatment of quality of information in sensing applications. For example, Hossain *et al*[14] describe quantification of similar quality of information attributes to those discussed by Bisdikian,[4] but in relation to face recognition. The advantage we offer in this latest work is in the synthesis of formal models which key into the total sensing and response problem at whatever levels of abstraction are found to be necessary. While there have been a number of treatments of low level measures of accuracy and related attributes, the stochastic modeling approach we have initiated will provide components which support links between the physics of measurement and the fulfillment of complex demands by command in military applications.

The advantageous novelty we offer springs from the proposal of a means for capturing the practicalities of military operations in flexible, formal stochastic models which reproduce the necessary degrees of freedom in a manner which enables the calculation of meaningful measures of quality of information. We have constructed a simple candidate quality of information metric describing the accuracy of an update from a sensor network, and shown how that quality decays over time. The research directions from this point include experimentation with different types of sensor, multiple sensor deployments in a given location, network protocol design, and mission characteristics.

ACKNOWLEDGMENTS

The authors would like to thank their colleagues in the International Technology Alliance for creating the energetic and open collaborative environment which fostered this work.

REFERENCES

[1] S. Ehikioya, “A characterization of information quality using fuzzy logic,” in *18th Int’l Conf. of the North American Fuzzy Information Processing Society,(NAFIPS 1999)*, pp. 635 – 639, June 10-12, 1999.

[2] T. Dasu and T. Johnson, *Exploratory Data Miningand Data Cleaning*, Wiley Series in Probability and Statistics, Wiley Interscience, 2003.

[3] “The MIT total data quality management (TDQM) program.” http://web.mit.edu/tdqm/www/index.shtml.

[4] T. P. C. Bisdikian, R. Damarla and V. Thomas, “Quality of information in sensor networks,” in *1st Annual Conference of ITA (ACITA’07)*, Sept. 2007. http://www.usukita.org/?q=node/148.

[5] J. Hillston, *A Compositional Approach to Performance Modelling*, Cambridge University Press, 1996.

[6] C. Priami, “Stochastic pi-Calculus,” *The Computer Journal* **38**(7), pp. 578–589, 1995.

[7] S. Gilmore and J. Hillston, “The PEPA workbench: A tool to support a process algebra-based approach to performance modelling,” in *Computer Performance Evaluation*, pp. 353–368, 1994.

[8] W. J. Knottenbelt, “Generalised Markovian Analysis of Timed Transition Systems,” Master’s thesis, Department of Computer Science, University of Cape Town, June 1996.

[9] M. Z. Kwiatkowska, G. Norman, and D. Parker, “PRISM: Probabilistic symbolic model checker,” in *Computer Performance Evaluation / TOOLS*, pp. 200–204, 2002.

[10] L. Cardelli and A. D. Gordon, “Mobile ambients,” in *Foundations of Software Science and Computation Structures: First International Conference, FOSSACS ’98*, Springer-Verlag, Berlin Germany, 1998.

[11] T. Osogami and M. Harchol-Balter, “Closed form solutions for mapping general distributions to quasi-minimal ph distributions,” *Perform. Eval.* **63**(6), pp. 524–552, 2006.

[12] P. G. Harrison, “Reversed processes, product forms and a non-product form,” *Linear Algebra and Its Applications* **386**, pp. 359–381, July 2004.

[13] I. Mitrani, *Probabilistic modelling*, Cambridge University Press, New York, NY, USA, 1998.

[14] M. A. Hossain, P. K. Atrey, and A. El Saddik, “Modeling quality of information in multi-sensor surveillance systems,” *Data Engineering Workshop, 2007 IEEE 23rd International Conference on* , pp. 11–18, 17-20 April 2007.