12th ICCRTS
"Adapting C2 to the 21st Century"

Performance Assessment of the C2ISR Enterprise

Topics:
S2 Metrics and Assessment,
Modeling and Simulation,
Networks and Networking

Dr. Michael B. Hurley (POC)
MIT Lincoln Laboratory
244 Wood Street
Lexington, MA 02420-9185
(781) 981-5049
hurley@ll.mit.edu

Mr. Peter Jones
MIT Lincoln Laboratory
244 Wood Street
Lexington, MA 02420-9185
(781) 981-0502
jonep@ll.mit.edu

Abstract

The substantial investment in sensor technology by the Armed Forces, coupled with the emergence of network services to connect consumers and producers, has resulted in a data glut. The military's command, control, intelligence, surveillance, and reconnaissance (C2ISR) community have rushed to acquire standards-based information enterprises to more efficiently manage this data, extract information, and make appropriate decisions to achieve mission objectives. These acquisition processes are occurring with very little theoretical or practical understanding of how to assess the performance of these large distributed enterprises. This paper presents the results of a study that developed a conceptual model and an analytical framework for the assessment of the C2ISR enterprise, with quantitative measures of performance and effectiveness derived from probability theory, information theory and utility theory. A simple simulation of a multi-sensor enterprise that develops a common operational picture was written to demonstrate the value of information theoretic measures for performance assessment. The simulation was used to assess different data communications architectures by evaluating the correctness, confidence, and consistency of the pictures that the architectures produce. Although the analysis was limited to assessing different communications architectures, the framework provides unified measures that can support trade-off studies between any set of disparate components in the C2ISR enterprise.

This work was supported by the Department of the Navy, Office of Naval Research (ONR) under Air Force Contract FA8721-05-C-0002. Opinions, interpretations, conclusions, and recommendations are those of the authors and are not necessarily endorsed by the United States Navy or the United States Air Force.

Motivation of study

The Integrated Sensing and Decision Support (ISDS) Group at MIT Lincoln Laboratory was created in 2004 to rapidly develop expertise in a new growth area for the laboratory. For decades, the dominant capability gap for the Department of Defense (DoD) has been insufficient information for decision-making. A multitude of sensor development projects have been funded through the decades to close the gap and Lincoln Laboratory has participated in a significant number of them. With the turn of the century, the new problem is data overload: how to manage the large volumes of raw data, convert the data to meaningful information, and use it to make the right decisions. The mission of the ISDS group is to find the solutions to this new high-priority problem.

The ISDS group began working with the Distributed Common Ground System – Navy (DCGS-N) Program Office in 2005 as the architect for the Navy's portion of the Distributed Common Ground System (DCGS). The goal of the DCGS program is to break down the stovepipes of the legacy Intelligence, Surveillance, and Reconnaissance (ISR) systems so that all pertinent information is always available to all ISR analysts. The DCGS program has adopted a Service-Oriented Architecture (SOA) for the technological foundation of this new information enterprise system. The significant gap that the DCGS acquisition program faces is the lack of a foundation to guide the performance evaluation of the SOA based ISR enterprise.

Working to narrow this gap, the ISDS group at MIT Lincoln Laboratory has completed a study for the Office of Naval Research to search for a foundation for the performance assessment of the ISR enterprise. The study assessed the state of the art in the evaluation of enterprise systems and examined other remotely connected fields like situation awareness and command and control for other frameworks that might be appropriate to the ISR enterprise. A suitable framework built upon a strong mathematical foundation was not found, and so one was constructed. The mathematical foundation was based upon probability theory, information theory and utility theory and was used to identify Figures of Merit (FoM), Measures of Effectiveness (MoE), and Measures of Performance (MoP) for enterprise architectures. In the final stage of the study a simulation was developed to demonstrate the utility of the framework's MOEs and MOPs for a C2ISR enterprise.

The study was originally a search for a single figure of merit for the ISR enterprise and was inspired by the radar and sonar equations. These equations provide measures that engineers use to assess the relative impact of system design decisions on overall system performance. An equivalent performance measure for the ISR enterprise would be a powerful tool for engineers to use for equivalent trade studies. The challenge that the study faced was how to move from a set of equations that characterize the sensitivity of single sensors to a set of equations that characterize an enterprise composed of multiple sensors, analysis systems, command and control elements, and actors.

The framework that was eventually developed during the course of the study showed that no single, objective figure of merit could be defined for the ISR enterprise, but that an

objective methodology could be defined for generating figures of merit for different ISR enterprise architectures. The framework also showed that figures of merit were best estimated for the combined C2ISR enterprise and not the ISR or C2 enterprise in isolation.

Literature Review – Assessment Techniques for Information Enterprises

The hope for the literature search was the discovery of an existing conceptual model with a strong mathematical foundation that would be suitable for the evaluation of the ISR enterprise. The study team examined a number of documents across the Department of Defense and university research and development community.

In a final report on the analysis of the Silent Hammer Limited Objective Experiment[1] a methodology is described for evaluating a multi-INT ISR network. The methodology included FoMs, MoEs, and MoPs that were used to assess the value of the future SSGN submarine to various Navy missions. The evaluation was analytical and methodical with good connections between the FoMs, MoEs, and MoPs, but the methodology did not have the strong mathematical foundation that is needed for the evaluation of the ISR enterprise.

The Office of Force Transformation's Network Centric Operations Conceptual Framework[2] (NCOCF) was examined for its applicability to ISR enterprises, especially its focus on network centric operations. Although the conceptual framework is well developed with proposals for metrics and measures, it contains no strong mathematical foundation that unifies the measures. The NCOCF led to other papers describing how to measure the quality of information and its impact on shared awareness[3] and a paper on complexity theory and network centric warfare[4]. These papers provided mathematical formulae for assessing networks based upon information theory, estimation theory, and complexity theory. The theories that these authors had selected were very strong candidates for the sought-for mathematical foundation. The relationship between the mathematical measurements and conceptual model was not described to a sufficient depth so that the connection between a mathematical foundation and conceptual model could be understood.

Kadambe and Daniell's paper[5] on the performance analysis of distributed sensor networks was examined for potential applicability to the evaluation of the multi-INT ISR enterprise. This paper proposed that the Euclidean distance, Shannon information, mutual information, and the symmetric Kullback-Leibler distance could assess the performance of distributed sensor networks. Some of the measures looked to be appropriate selections for evaluating an ISR enterprise, but a unifying mathematical foundation that tied these measures together into a unified whole was missing from the paper. Additional work would also be needed to extend the metrics to a mathematical founded conceptual model for the ISR enterprise.

Mica Endsley's model of situation assessment[6, 7] was also investigated. This model is well known and provides a very good conceptual model for understanding systems that assesses situations. It has proven to be very useful for the human factors community in assessing human operators' awareness of and ability to correctly assess situations. The model does not have a mathematical framework, although statistical techniques are often used to evaluate situation assessment systems.

Mahoney, Laskey, Wright, and Ng's paper[8] on situation awareness proposes specific mathematical formulae for evaluating systems like those in the Endsley model. This paper suggested that probability density functions (PDF) be adopted as a mathematical representation for groups, units, sites, and activities. The systems described in this paper are constructed with Bayesian networks and develop probabilistic estimates of situations. The PDFs can be used to evaluate the accuracy of the situation estimate, both at the global and local levels of awareness. Although the paper described a mathematical framework for the systems, the proposed scoring functions for evaluating the system looked to be ad hoc with little theoretical justification for its formulation.

Finally, a number of different papers on the assessment of information quality were reviewed. Knight and Burn's paper provides a good overview of the information quality literature[9]. So far, this community has been more interested in defining information quality, often compiling lists of up to 20 different attributes of information quality, rather than in developing a model of how information is created and used.

In summary, the literature search identified many conceptual models that many common attributes and that could be relevant in a conceptual model of the ISR enterprise, but no conceptual model was found that had a strong analytical foundation and a single Figure of Merit. The team chose to develop a conceptual model with a strong mathematical foundation in the next phase of the study with the results of the literature search guiding the shape of the new conceptual model.

The Conceptual Model

The study team decided that a new mathematically based conceptual model was needed to adequately describe the ISR enterprise. The model was developed with three levels of complexity, with appropriate mathematical theories identified for the performance evaluation of any level of the model. All three levels of the model were developed in an integrated, iterative fashion, but for expositional clarity each level is detailed separately below.

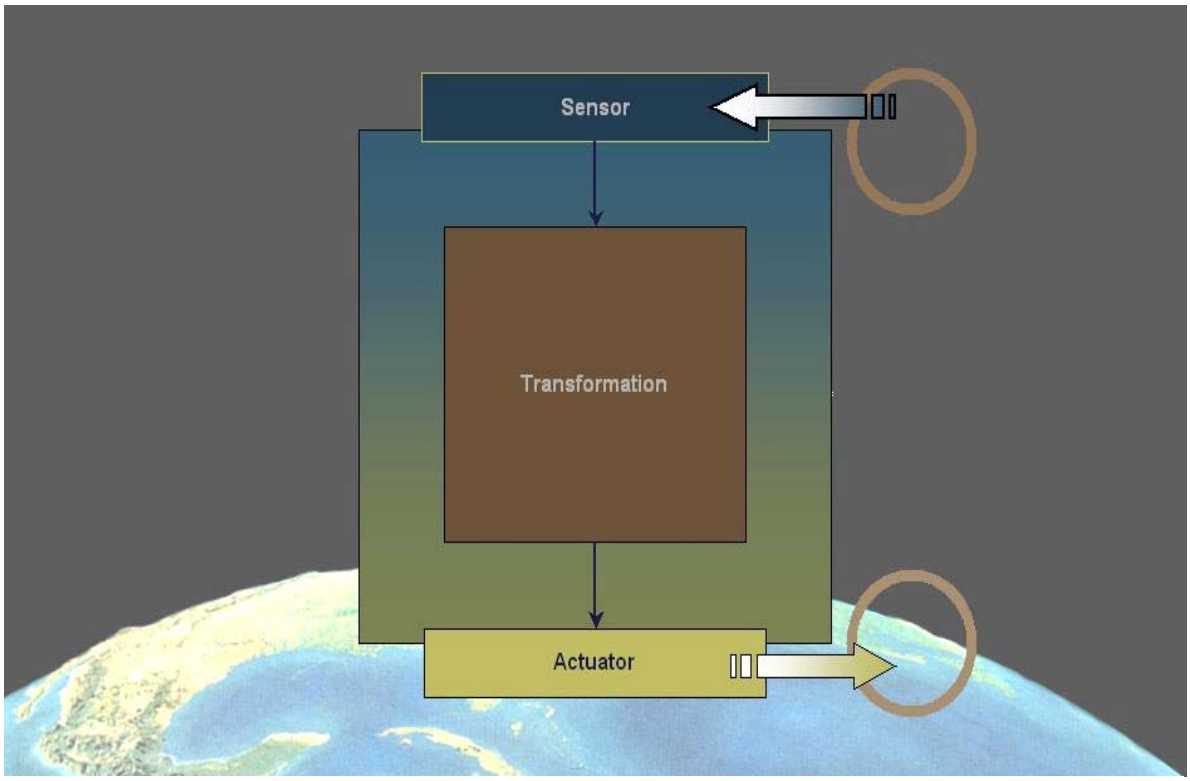

Figure 1 The first-order model.

The First-order Model – the Black Box

The simplest model of a decision system is of a black box existing in an environment; it has sensors that measure aspects of (or perceive) the environment and actuators that can change the environment. The nature of the measurements that the decision system makes may induce it to use its actuators to alter the environment. This simple model is shown in Figure 1. This decision system forms the fundamental unit for the more complex decision models that follow.

The minimum mathematical description for this model is a mapping from sensor measurements to actuator events:

$$A = T(S, i), \tag{1}$$

where A is the chosen action, selected via the transformation function T, the sensed data S, and any internal or prior data, i, inside the black box. This forms the mathematical basis for all the components in the decision models that follow, where, generically, S is the input into the component and A is the output.

Almost all models found during the literature search are extensions of this basic model, including the Endsley model, the Office of Force Transformation's Network Centric Operations Conceptual Framework, and (in retrospect) Boyd's OODA loop[10]. This conceptual model explicitly interprets the C2ISR enterprise as a system that transforms data and information into action.

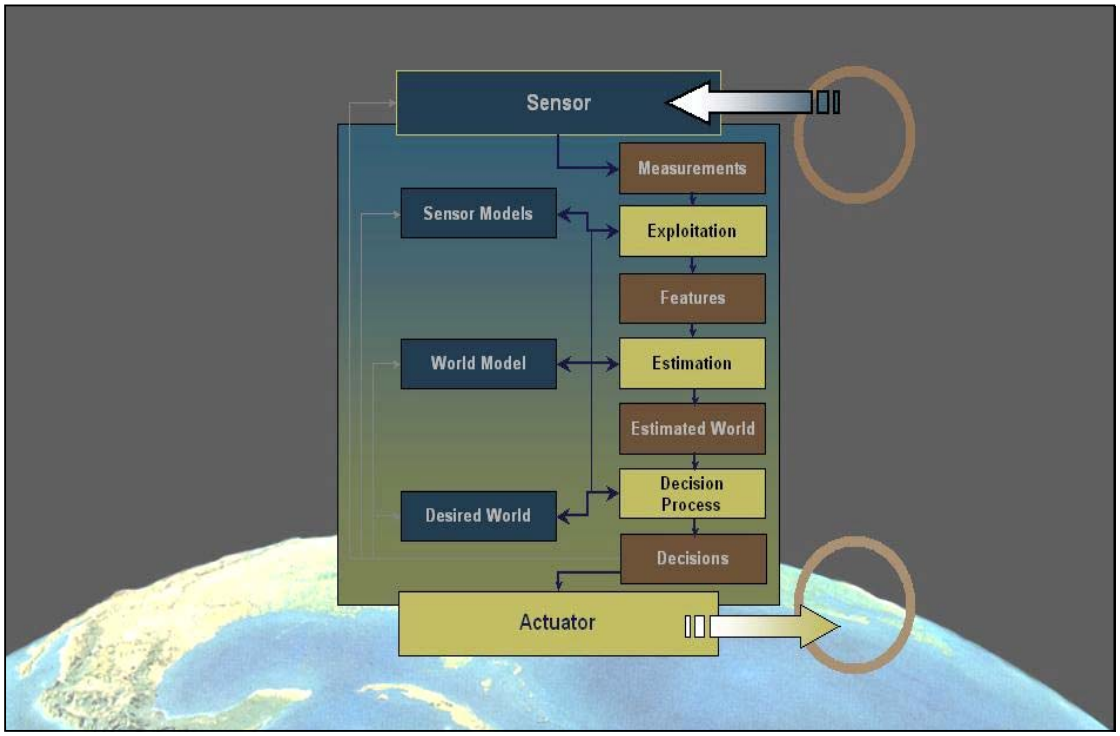

Figure 2 The second-order model.

The Second-Order Model – the Simple Decision System

As model development proceeded, it was eventually realized that the ISR enterprise needed to be modeled as a part of a larger decision system that encompassed both C2 and ISR enterprises. Clearly, the evaluation of the ISR enterprise would be closely tied to the evaluation of a complete decision system.

The second-order model defines the contents of the black box of the first-order model and describes the data flow between the sensors and the actuators: the decision system has sensors that sense the environment and report measurements, an exploitation process that converts the measurements into features, an estimation process that collects the features and constructs an estimate of the world, a decision process that uses the estimate of the world to decide upon the appropriate course of action, and actuators that act upon the decision to change the state of the world, including the decision system's internal models. This flow is shown in Figure 2. Each process can be mathematically formulated as a series of transformations like that of Equation (1). The figure also shows a set of internal models that the processes in the chain may use to convert input data into output data. These are equivalent to the variable, i of Equation (1), and represent internal or prior data and are shown in the figure as sensor, world, and desired world internal models. The internal data may be accessible to any of the processing components of the second order model because there are no communications costs for the closely connected components. A learning mechanism is included in the conceptual model in Figure 2 by allowing for the decision process to change the internal models based upon the world state estimate.

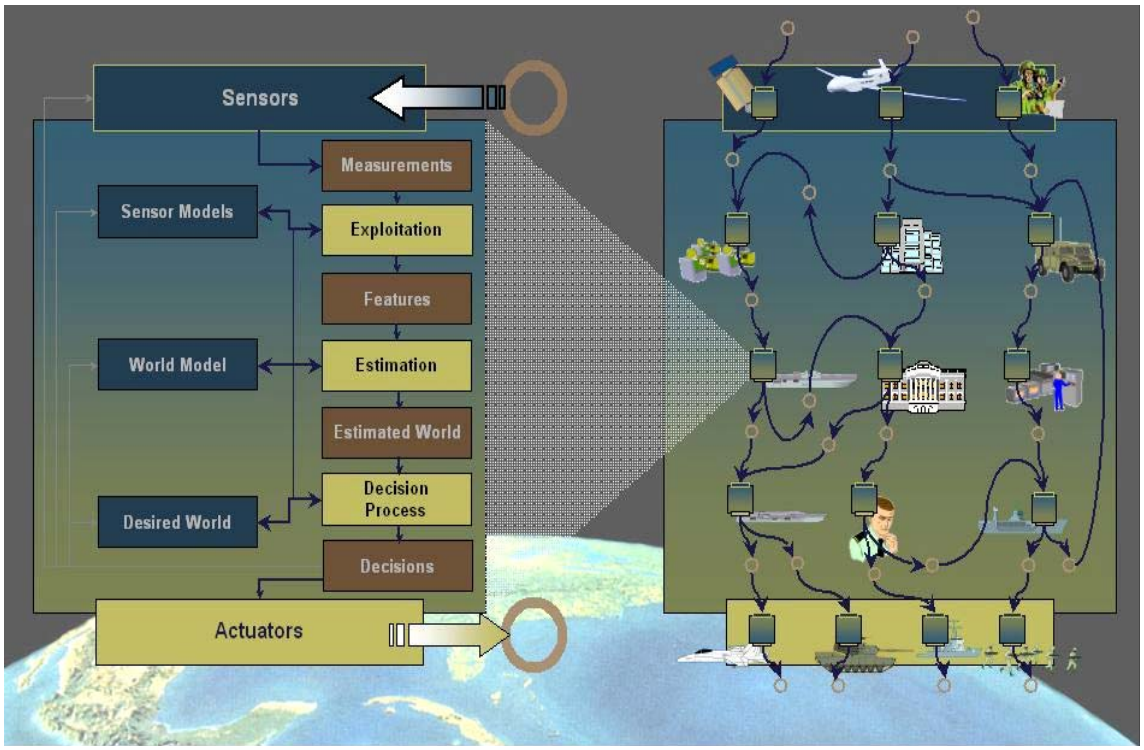

Figure 3 The third-order model.

The second-order model begins to diverge from the models found during the literature search, although many aspects are similar to Endsley's model, the Office of Force Transformation's Network Centric Operations Conceptual Framework, and (in retrospect) Boyd's OODA loop. The differences are partly by the different focuses of the model developers. The focus of this study was the development of an enterprise assessment methodology with a mathematically based conceptual model necessary to drive the development of this methodology.

The Third-order Model – the Enterprise Decision System

The second-order model was known to not be complex enough to adequately describe the C2ISR enterprise, which is composed of multiple, loosely coupled, spatially distributed components that interact to achieve a common goal. The final extension to the model was to link multiple second-order decision systems into a single distributed system, as shown in Figure 3. Components in the third-order model are assigned primary responsibility for performing the work of individual processes in the second-order model. In the C2ISR enterprise, ISR units are assigned primary responsibility for converting sensor measurements into features and C2 units are assigned primary responsibility for organizational decisions. The responsibility for the process of constructing the world state estimate has often been split between ISR units and C2 units within the DoD community. Unlike the second-order model, the assumption for the third-order model is that there are costs associated with communications between components. Because of the communications limitations, the sensors, world, and desired world models do not exist in the third-order model, other than through their containment in second-order components.

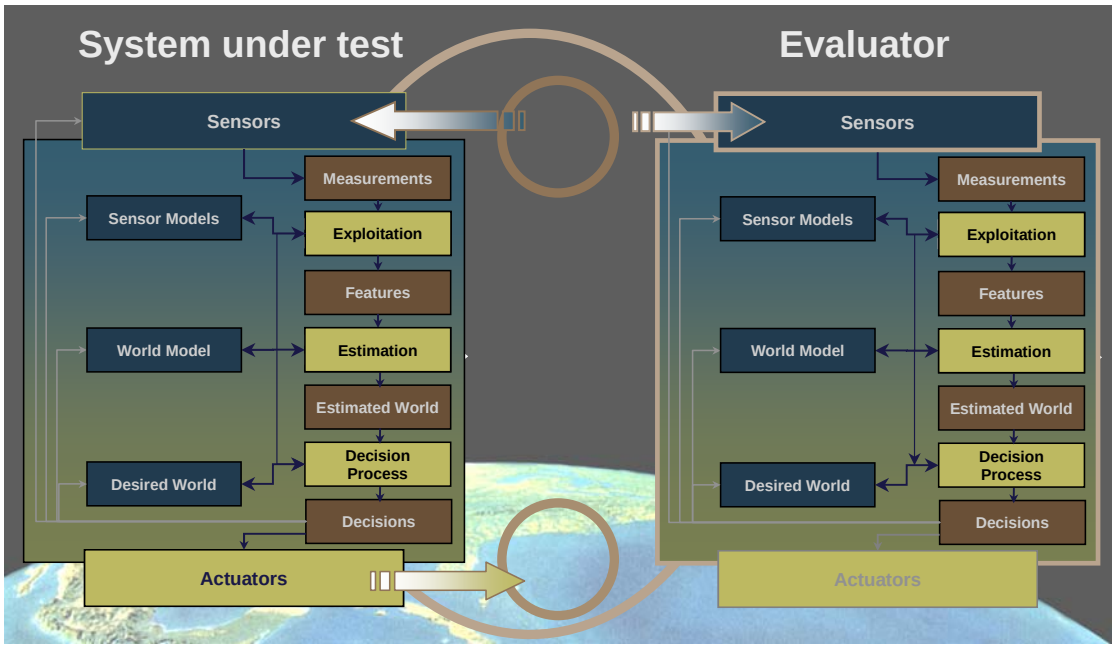

Figure 4 The evaluator as a decision system.

The Evaluator as a Decision System

Because the goal of the study was to develop a performance evaluation methodology for the C2ISR enterprise, the conceptual model was extended to include an evaluator who performs performance assessments on other decision systems, as shown Figure 4. The evaluator senses the environment, converts sensor measurements into features, and the features into a world state estimate. The evaluator uses the world state estimate to determine the value of the decision system under test. The evaluator's actuators are used to control and configure the environment in preparation for tests, and to report the results of the tests. In most cases, the evaluator will be able to better assess performance if it can sense more of the environment than the decision system and if it can also sense the internal states of the decision system. An alternative for cases where the internal states of the decision system are inaccessible to the evaluator is to have cooperative decision systems write internal state data into the environment through its actuators.

An insight from this model is that evaluators do not necessarily have to know the explicit details of the internal processes within a decision system, but can model those internal operations with a methodology of their own choice as long as it captures the aspects that are important to the evaluator. Figure 4 shows an evaluator assessing a second-order decision system, but evaluators can evaluate any order of the model. The evaluator can even model higher-order systems as lower-order systems to simplify their evaluation process. A model of the evaluator provides insight into how evaluators assess decision systems and how the mathematical foundations influence the kinds of assessments that evaluators can make.

Mathematical foundations

A number of different mathematical theories were considered for potential foundations of the conceptual models. The development of the conceptual model and the selection of a mathematical foundation was an iterative process with both tasks influencing the other. The mathematical foundation was eventually built up from probability theory, information theory, and utility theory. Of the three theories, probability theory is a foundation of the other two and probability density functions (PDFs) are the fundamental mathematical object in the model. This agrees with the work of Mahoney et al, who also used PDFs as the fundamental object in their models.

Probability Theory

The mathematical foundation was integrated with the conceptual model by assuming that an evaluator can model a decision system and its environment as a state machine. The environment and decision system can be separated into two orthogonal state spaces. Both the environment and decision system are in one unique state at any given moment, but no decision system has full knowledge of the true states. This uncertainty can be expressed with probability density functions. The evaluator views the series of processes in the second- and third-order decision models as performing conversions from one probabilistic state representation to another, whether or not the decision system uses the same probabilistic representations as the evaluator. To clarify this point, the sensors measure aspects of the environment and report measurements that an evaluator interprets as probabilistic estimates of the real environmental sub-states that are being measured. The exploitation process converts the measurement PDFs to feature PDFs, the estimation step generates a PDF over all interesting world states, and the decision process uses the world state PDF to select a decision.

The consistency between two PDFs, when the true value of a measurement, feature, or world state space or subspace, x, is not known can be measured with the probabilistic function,

$$D_B = \sum_x P_E(x) P_D(x), \tag{2}$$

where D_B is the numerical correctness measure, $P_E(x)$ is the first PDF, possibly estimated by an evaluator, and $P_D(x)$ is the other PDF, possibly estimated by a decision system. The consistency measure can be used to compare the any two PDFs with the same spaces from different second-order decision systems in the third-order model.

If the evaluator has a significantly more accurate estimate of the true state x, then Equation (2) is a reasonable measure of the correctness, or accuracy of the decision system's estimate. If the evaluator knows the true value, x_t, with absolute certainty, then $P_E(x)$ is a delta function, $\delta(x - x_t)$, and the measure becomes

$$D_B = P_D(x_t). \tag{3}$$

In general, a good decision system's PDFs will agree more closely to the evaluator's estimate of the truth than would the PDFs of a bad decision system.

Information Theory

The chain of processes in the second-order model forms a chain of information channels, motivating the selection of information theory as the second of the mathematical foundations. The third-order model, in general, is a network of information channels. Both the second-and third- order models are amenable to analysis with information theory.

Shannon information, defined in Equation (4), provides a measure of the uncertainty (or entropy) of a PDF:

$$H(x) = -\sum_{x} P(x)\ln(P(x)). \tag{4}$$

Smaller entropies correspond to more confident estimates of the state. This measure is only related to precision and not accuracy; it does not reflect whether an estimate is correct. A PDF can be completely inaccurate but still have a small value for its entropy. The inversion of Shannon information can be considered to be a measure of confidence.

The relative entropy, or Kullback-Leibler distance, is another measure of the similarity of two PDFs:

$$D(P_E \parallel P_D) = \sum_{x} P_E(x)\ln(P_E(x)) - \sum_{x} P_E(x)\ln(P_D(x)). \tag{5}$$

Relative entropy is always non-negative, and is zero if and only if $P_E(x) = P_D(x)$. The convention is that $0\ln(0) = 0$ and $-P_E(x)\ln(0) = \infty$. This measure is not a metric distance because it is not symmetric; the exchange of the two PDFs does not produce the same value (i.e. $D(P_E \parallel P_D) \neq D(P_D \parallel P_E)$). A symmetric version of the Kullback-Leibler distance, D_S, is often used to avoid specifying the preferred PDF and ensuring that the same value is produced independent of the PDF order in the function:

$$D_S(P_E \parallel P_D) = D(P_E \parallel P_D) + D(P_D \parallel P_E). \tag{6}$$

The Kullback-Leibler statistical distance and its symmetric version can be used to measure the consistency between PDFs at different locations in an enterprise or decision system. As with the probabilistic correctness measure, the two PDFs must be defined on the same sub-spaces.

Because the processes in a decision system and an enterprise are communications channels, the information loss that occurs due to inefficient transmission can be measured. Mutual information provides a measure of the common information, or the consistency of information, on both sides of the channel,

$$I(X;Y) = -\sum_{x} P(x)\ln(P(x)) + \sum_{x,y} P(x,y)\ln(P(x \mid y)). \tag{7}$$

An evaluator can use this measure for series of channels to identify collective inefficiencies. This measure requires that the evaluator construct a PDF, $P(x,y)$, over the combined space of both x and y that are the subspaces at the beginning and end of the

channel. The conditional PDF, $P(x \mid y)$ is also required, but can be derived from the overall PDF, $P(x, y)$, through standard probabilistic integrals. The most difficult part of calculating this measure is collecting measurements from both spaces and generating a PDF estimate over the combined space. It may be noted that the two spaces do not have to be defined over the same sub-spaces.

Utility Theory

Probability theory and information theory provide measures of correctness, confidence, and consistency, but do not provide a single unified figure of merit. An evaluator is more often interested in the value of a decision system for achieving the goals of the evaluator. Utility theory provides measures of value and is the final theory that was selected to form the mathematical foundation for the model. It can be used to estimate the overall cost or benefit of a decision system from the perspective of an evaluator. The equivalent financial measure would be total lifecycle cost.

Utility theory requires that an evaluator place a value on all states of the environment. This means that there cannot be a single unified figure of merit for the C2ISR enterprise because different evaluators will assign different values to world states – in other words, utility theory provides measures that are relative. Only if evaluators can agree on common value assignments can they can agree on a common figure of merit. However, the theory does provide a uniform methodology for the evaluation of decision systems,

An evaluator can conduct repeated tests on a decision system to construct a probabilistic estimate of the environmental states that a decision system drives toward. The combination of the evaluator's value assignments with the probabilistic estimate produces an average utility value for a given decision system. In all meaningful evaluations, either multiple decision systems are compared with respect to each other or a decision system is compared to an environment lacking a decision system (a null decision system for those wanting a unified description of the comparative process as always being between two decision systems). An evaluator uses the average utility values to determine which decision system is most valuable.

The average cost $\langle C \rangle$ for a decision system can be estimated with

$$\langle C \rangle = \frac{\int_{t_{start}}^{t_{end}} \sum_{E} C(E) P(E \mid x, i, t) dt}{t_{end} - t_{start}}, \tag{8}$$

where t is time, $C(E)$ is the cost function for all environmental states, E, and $P(E \mid x, i, t)$ is the conditional PDF that the environment is in state E, given an evaluator's measurements, x, prior information, i, and valid a time t. The average cost is a sum over all environmental states, and averaged over the duration of the test, $t_{end} - t_{start}$. Other variations to the equation are possible, depending on the form of the information that an evaluator wants to use in generating the average.

Retrospective Literature Review

In a retrospective literature review conducted in preparation for this paper, additional models and mathematical foundations were examined, with a greater emphasis given to the C2 community. The NATO C2 conceptual model[11] was missed in the original review, partly because it was published and released during the original literature review. At least a partial mapping can be made between the NATO C2 conceptual model and the conceptual model that is presented here. Because the NATO C2 conceptual model is a comprehensive model, it has not yet been completed and a mathematical foundation still remains to be developed to support the long list of variables and metrics that are already identified in the model.

Papers by S. L. Forsythe, et al.[12], [13] have a similar conceptual feel to the models of this paper. These researchers propose a multi-level model with nested C2 loops that is similar in spirit to the model presented here. Their C2 loop is an aggregation of a number of other C2 models, including the OODA loop, MAAPER, models by David Noble and Jay Bayne, and a C2 Joint integrating Concept (JIC) document.

A paper by N. Smith and T. Clark[14] examines multi-attributive utility theory (MAUT) and Value Focused Thinking (VFT) as potential mathematical foundations for evaluating C2 systems and determines that VFT is the right way to assess complex domains like net-centric systems. MAUT is rejected primarily because of the strong interdependence between model parameters. Bayesian networks and influence diagrams are adopted as mathematical tools to track the dependencies between the important parameters of the C2 system under evaluation.

A number of papers by Perry et al and Moffat were examined during the original literature search and additional papers[15] continue to be published, especially with Perry and Moffat[16] collaborating on the more recent papers. These papers are very close to the mathematical foundations proposed here with a number of similarities. For example: Information theory is a principal component of the mathematical foundation, which includes probability theory by default. The entire decision system must be included in the evaluation process. Complex organizations are composed of hierarchical layers that make evaluation difficult.

There are differences between this paper and the publications of Perry and Moffat, but the similarity between the multiple efforts is very striking. The differences between the teams can be explained by the immaturity of effort to apply information theory to C4ISR models and by the amount of work still needed to fully integrate mathematical foundations into models of complex C2ISR systems.

Simulation of a Multi-INT Multi-Sensor Enterprise

The figure of merit, measures of effectiveness, and measures of performance as provided by utility theory, information theory, and probability theory can assess the C2ISR

enterprise. If the C2ISR enterprise is decomposed into a C2 enterprise and an ISR enterprise, a utility assessment of the C2 enterprise might be able to proceed with a simplified model of the ISR enterprise, but a utility assessment of the ISR enterprise is nearly impossible because the organizational goals are in the domain of the C2 enterprise. A utility assessment for the ISR enterprise might be accomplished if the C2 enterprise is considered to be part of the environment and not part of the decision system. Assessment of the ISR enterprise would then be based upon the goals of the ISR enterprise and not the C2ISR enterprise. This utility assessment is very likely not the kind that high-ranking military officers would like to see within their organization.

Because the original focus of the study was limited to the ISR enterprise and it was learned rather late in the study that the C2 enterprise must be included to perform utility assessments, the simulation that was developed for demonstration of an evaluation of the C2ISR enterprise was restricted to the ISR enterprise. Although a utility assessment could not be made, the evaluation of the ISR enterprise with the measures of effectiveness and performance could still be demonstrated in the absence of knowledge about the C2 enterprise. The simulation used probabilistic and information theoretic measures of correctness, confidence and consistency to evaluate the simulated enterprise.

Simulation Description

The simulation was chosen to reflect a topic of interest in the C2ISR community: the creation of a common operation picture (COP) by a set of heterogeneous loosely connected decision systems. The COP is equivalent to the world state estimate that is generated by the estimation process in the second-order model. Most of the information systems that currently generate COPs are C2 systems and not ISR systems, illustrating that the boundary between the two communities is fuzzy, especially in light of the conceptual model.

The enterprise consisted of five sensors, each with a unique sensing modality, which traveled through the grid world observing different portions of the world. The sensor capabilities are shown in Table 1 and Table 2.

	Sensor A	Sensor B	Sensor C	Sensor D	Sensor E
O	1	3	5	1	1
□	1	3	1	5	1
Δ	1	3	1	1	5
None	0	0	0	0	0

Table 1 Sensor detection strengths as a function of target class

	Sensor A	Sensor B	Sensor C	Sensor D	Sensor E

Extent	5	3	1	1	1

Table 2 Sensor observation extent (edge length of observable grid)

The simulated environment consisted of a 5 x 5 grid world. Each grid location contained at most one object of interest. In reality, these could be vehicle types or combatants, but in the simulation, objects were represented generically as circles, squares, and triangles. Thus, each grid location contained one of these three target classes or nothing at all.

As an example of how to interpret the tables, if sensor B were located at a specific cell in the 5 × 5 grid, it would be able to observe all nine cells in the 3 × 3 sub-grid centered on its location. For each observed cell that contains a circle, square, or triangle the resultant measurement would be $3+\eta$, where η is an additive noise term. If the cell does not contain any of the objects, the resultant measurement would be $0+\eta$. (All noise terms are drawn from zero mean, unit standard deviation Gaussian distributions). It is significant to note that sensor B has no ability to discriminate between a circle, square, or triangle; all return statistically identical measurements. This means that without data sharing from other sensors in the enterprise it is impossible for sensor B to discriminate between all the objects of interest. The same is true of each of the other four sensors.

In the experiment, five different enterprise architectures (described in the paragraphs below) were run through a series of 300 Monte Carlo trials each. The information theoretic measures of entropy, Kullback-Leibler (KL) distance and Bayesian agreement with truth (Bayesian Distance) were collected and aggregated across the trials. Figures 5–9 demonstrate the power of these metrics in evaluating the performance of the various architectures.

Evaluation of Communications Architectures for a Multi-Sensor Enterprise

The first architecture that was examined was termed the “unlimited communication” architecture. In this architecture, there were no limitations on communication bandwidth and so the ability of the enterprise to share and assimilate knowledge was limited solely by the rate of data acquisition. The sensors fused all measurements from all sensors into their local world models with Bayesian logic. This idealized architecture provides an upper bound on the levels of expectable performance. The statistical results are shown in Figure 5, with each sensor represented as a colored line: blue is sensor A, green is sensor B, etc. The plots demonstrate a rapid convergence to absolute certainty (measured by Shannon entropy), complete consistency (measured by the symmetric Kullback-Leibler distance), and perfect correctness (measured by Bayesian agreement with truth). The small amount of initial inconsistency is due to latency effects, and the convergence rates of the other two metrics are due to the noise characteristics of the environment.

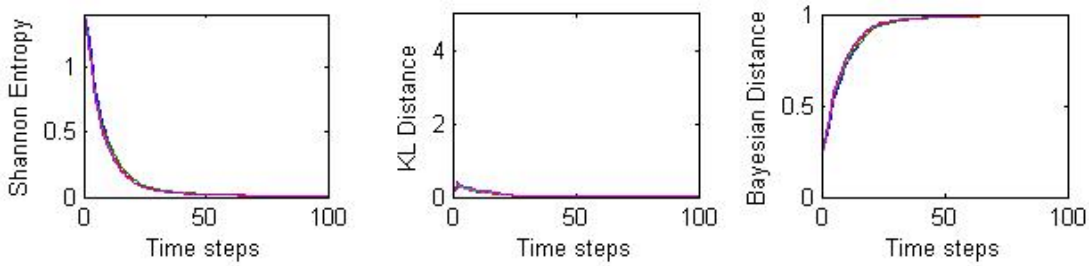

Figure 5 Average Shannon entropy, KL distance, and Bayesian agreement with truth for each of the five sensors when using the "unlimited communication" architecture.

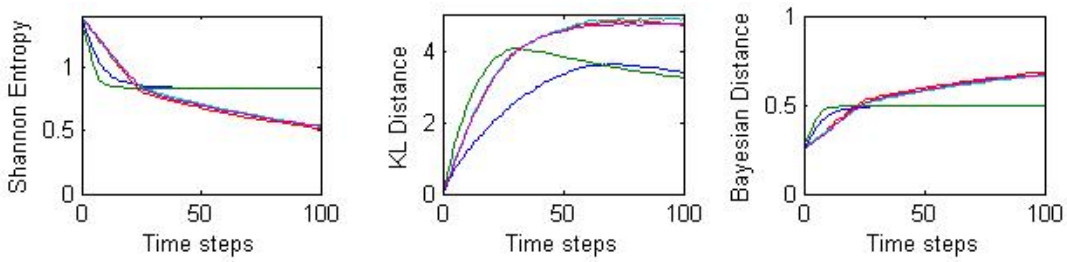

Figure 6 Average Shannon entropy, KL distance, and Bayesian agreement with truth for each of the five sensors when using the "no communication" architecture.

In the second architecture, called the "no communication" architecture, there was no communication between any of the decision systems. This worst-case scenario was used to establish a lower baseline on the potential performance of the other architectures. In this architecture the sensors updated their local world models using Bayesian logic, but did not communicate any gained knowledge with any other sensors in the enterprise. The results, shown in Figure 6, are that the world models of the various sensors do not ever achieve a high degree of certainty, consistency or correctness.

The final three architectures that were investigated all assumed a uniform, fixed bandwidth constraint on node-to-node communication and were destined to be somewhere between the first two architectures in terms of performance. In the first architecture, called the "blind push" architecture, all sensors attempted to push all measurements to all other sensors. Due to the communication constraints, this resulted in first-in, first-out (FIFO) queue of measurements at each node. The performance results of this architecture are shown in Figure 7. Because of the absence of a prioritization mechanism, important measurements got jammed in the communications bottlenecks and all the performance measures suffered accordingly. While not as poor as the "no communication" case, there is significant performance degradation from the optimal "unlimited communication" case.

The second fixed bandwidth architecture was termed the "blind pull" architecture. In this architecture sensors were allowed to broadcast a request (which took a portion of the bandwidth) at each time step. These requests enabled the other sensor nodes to prioritize the measurements they pushed to that sensor node. However, because the broadcast request consumed some bandwidth, fewer measurements could be communicated. The results, shown in Figure 8, show that despite receiving fewer measurements from other

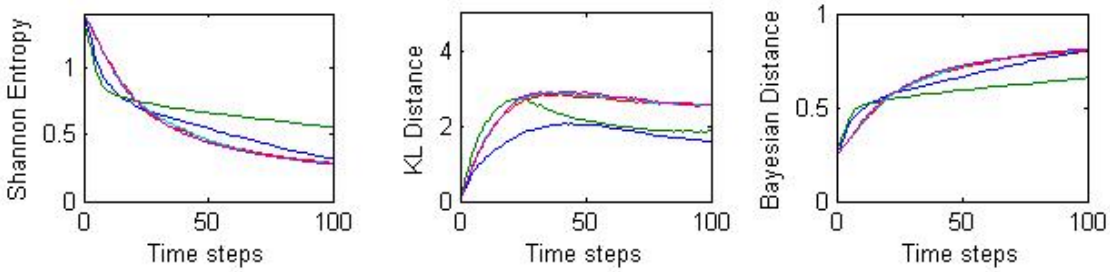

Figure 7 Average Shannon entropy, KL distance, and Bayesian agreement with truth for each of the five sensors when using the "blind push" architecture.

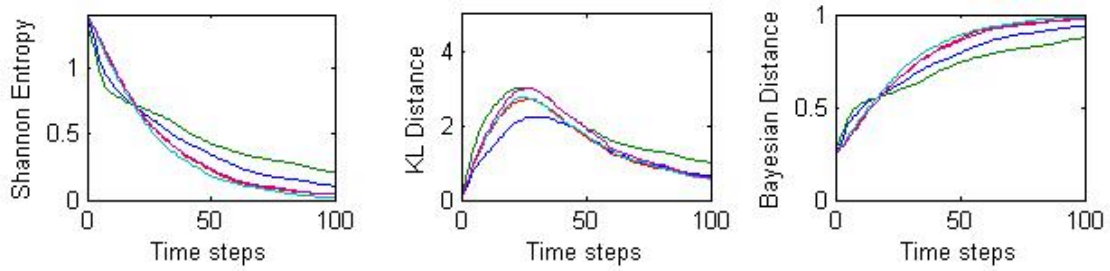

Figure 8 Average Shannon entropy, KL distance, and Bayesian agreement with truth for each of the five sensors when using the "blind pull" architecture.

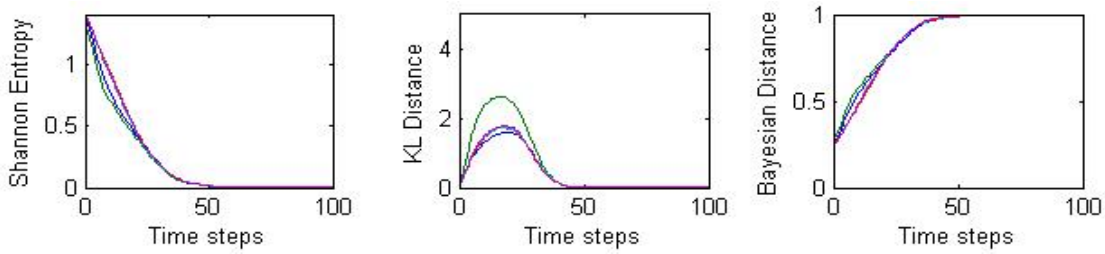

Figure 9 Average Shannon entropy, KL distance, and Bayesian agreement with truth for each of the five sensors when using the "informed pull" architecture.

sensors, the rudimentary prioritization mechanism improved enterprise performance over the "blind push" case.

The final architecture that was investigated was called the "informed pull" architecture. In this variant, each sensor node broadcast its position information at each time step. This can be thought of as metadata for that sensor's measurement at that time step. Other sensors could then request specific measurements, as informed by the metadata. Transmitting the metadata and data transfer requests consumed a portion of the bandwidth, with each message approximately equal to 60% of a sensor measurement report, resulting in fewer measurements being communicated than in either the "blind push" or the "blind pull" architecture. However, the metadata enabled the other sensors to direct requests for specific data products, creating a very efficient prioritization scheme. The performance results are shown in Figure 9. There was a dramatic increase in the ability of the enterprise to maintain consistency among the distributed nodes, to quickly converge to a high level of certainty, and to correctly classify all the objects of interest.

Some Implications of the Model

There are a number of implications that can immediately be identified with the conceptual model. More may be found in time, as the model is refined for the evaluation of specific configurations of the C2ISR enterprise and is expanded to explain more of the higher-level constructs found in the C2 models.

The first is that a utility-based evaluation of the ISR enterprise cannot be conducted without knowledge about the sensors, C2 enterprise, and actuators (such as weapons systems) that form a complete decision system.

Second, traditional figures of merit like the radar and sonar equations evaluate how well a sensor measures an environment. Most of the other traditional figures of merit evaluate how well an actuator changes an environment. Service-oriented architectures and the C2ISR enterprise need figures of merit that evaluate how well a system extracts information from measurements of the environment and decides what actions to take.

There has been much discussion in the DoD community on the value of a common operational picture (COP) and variants like a common relevant operational picture (CROP) and consistent operational picture (COP). The third-order model implies that, even with a COP or CROP, the C2ISR enterprise may still perform poorly because the goals of individual components may not be consistent with the goals of the overall organization. The C2 concept of commander's intent captures the need for goals to align across an organization for optimal performance.

The presentation of three levels of complexity for decision systems is not meant to imply that there are only three levels. In fact, specific models may require deeply nested decision systems of systems to accurately model complex enterprises. The components of a organizational decision system may be heterogeneous, with components best described as first-level systems, partial second-level systems, and more-complex-than-third-level systems. For example, humans perform massively parallel computations and often make conflicted decisions, which may be conjectured to indicate that there are different modules in the brain that generate decisions which a higher-order module has to arbitrate in order to resolve potential conflict.

The possibility for deeply nested decision systems in the enterprise means that some components in the systems are complete, second-order or higher decision systems. These systems are capable of many more functions than just their primary role in the enterprise and are making their own decisions and pursuing their own goals, which may be in conflict with the larger organization. The most capable components may even be able to assume an evaluator role and attempt to influence the performance of other components to achieve their own interests rather than the collective interests of the enterprise.

Summary

A six-month study, executed by the ISDS group at MIT Lincoln Laboratory for the Office of Naval Research, has developed a conceptual model to describe the C2ISR enterprise. The conceptual model has a strong mathematical foundation based upon probability theory, information theory, and utility theory. Probability theory provides the fundamental objects in the model: probability density functions. Information theory provides the fundamental concept of decision systems: information networks that distribute information through various representations of probability density functions. Utility theory provides the fundamental figure of merit: a probabilistic estimate of the total utility of a decision system to an evaluator.

The theories provide general-purpose equations that can be applied to specific instances of enterprise evaluation. The more basic equations from these theories have been presented in this paper. The power of the general-purpose equations is that they integrate the contributions of the heterogeneous components of decision systems into unified, quantifiable measures of information and utility, and are analogous to the way that the radar and sonar equations integrate the contributions of heterogeneous components of radar and sonar systems into unified quantifiable measures of the overall sensitivity of the sensors. The three foundational theories can provide additional analytical equations for more complex analyses of enterprise systems, which have not been presented here because of space.

Significant research remains to extend the results of this short study into full-scale evaluations of real enterprise systems. The study was originally focused on the ISR enterprise and on developing a model with a strong mathematical foundation. As the study has progressed, similar efforts have been found to be underway in the C2 community. These models concentrate on the complex human and social factors that can affect a C2 organization. A significant amount of research remains to extend the mathematically based model to a point that it can explain the social-science-dominated models of the C2 community. Anticipated benefits from an information theoretic foundation for C2 conceptual models would be that information flows are correctly captured, all necessary components are included, the large lists of attributes that affect C2 enterprise performance can be unified, and measures identified that accurately evaluate these attributes.

---

[1] P.A. Pomianowski, R. Delanoy, J. Kurz, and G. Condon, “Silent Hammer,” Lincoln Laboratory Journal, vol. 16, no. 2, 2006.

[2] J.J. Garstka, “Network Centric Operations Conceptual Framework: Version 1.0,” Office of Force Transformation, Washington, DC (November 2003).

[3] W.L. Perry, D. Signori, and J. Boon, *Exploring Information Superiority: A Methodology for Measuring the Quality of Information and its Impact on Shared Awareness*, Rand Corporation, CA, ISBN 0-8330-3489-8 (2003).

[4] J. Moffat, "Complexity Theory and Network Centric Warfare," DoD Command and Control Research Program, Washington, DC (2003).

[5] S. Kadambe and C. Daniell, "Theoretic Based Performance Analysis of Distributed Sensor Network," HRL Labs, Malibu, CA, Final Technical Report, AFRL-IF-RS-TR-2003-231 (1 October 2003).

[6] M.R. Endsley, "Toward a Theory of Situation Awareness in Dynamic Systems," Human Factors, 37(1), (1995), pp 32–64.

[7] M.R. Endsley, "Measurement of Situation Awareness in Dynamic Systems," Human Factors, (37) 1, (1995), pp 65–84.

[8] S.M. Mahoney, K.B. Laskey, E.J. Wright, and K.C. Ng, "Measuring Performance for Situation Assessment," *Proc. of 2000 MSS National Symposium on Sensor and Data Fusion*, San Antonio, TX (June 2000).

[9] S-A. Knight and J. Burn, "Developing a Framework for Assessing Information Quality on the World Wide Web," Informing Science Journal, Volume 8 (2005).

[10] John Boyd (Col.), "Patterns of Conflict," December 1986, http://www.d-n-i.net/boyd/pdf/poc.pdf.

[11] D. S. Alberts, et al.., "NATO Final Report SAS-050 "Exploring New Command and Control Concepts and Capabilities," www.dodccrp.org/files/**SAS-050**%20Final%20Report.pdf, January 2006.

[12] S.L. Forsythe, P.D. North, and V.B. Barnes, "Evaluation of Net-centric Command and Control via a Multi-resolution Modeling Evaluation Framework," *10th International Command and Control Research and Technology Symposium*, June 2005.

[13] S.L. Forsythe and P.D. North, "A Process Decomposition Approach for Evaluating Command and Control (C2) Functional Performance," *2006 Command and Control Research and Technology Symposium*, June 2006.

[14] N. Smith and T. Clark, "A Framework to Model and Measure System Effectiveness," 11th International Command and Control Research and Technology Symposium, 2006.

[15] W. Perry, R. Button, J. Bracken, T. Sullivan, and J. Mitchell. "Measures of Effectiveness for the Information-Age Navy: The Effects of Network-Centric Operations on Combat Outcomes." MR 1449, RAND, 2002 [http://www.rand.org/].

[16] W. Perry and J. Moffat. "Information Sharing among Military Headquarters: The Effects on Decisionmaking," MG-226, RAND, Santa Monica, CA, USA, 2004.