18th ICCRTS:

C2 in Underdeveloped, Degraded and Denied Operational Environments

"Evaluating the Value of Information in the Presence of High Uncertainty"

Data, Information, and Knowledge

Experimentation, Metrics, and Analysis

Collaboration, Shared Awareness, and Decision Making

STUDENTS: David Sidoti, Diego Fernando Martínez Ayala, Xu Han, Manisha Mishra, Sravanth Sankavaram
E-mail: {sidoti, dfm08004, xuh06002, manisha.mishra, mss11009}@engr.uconn.edu

Dr. Woosun An
E-mail: woa05001@engr.uconn.edu

Prof. David L. Kleinman
E-mail: dlkleinm@nps.edu

Prof. Krishna R. Pattipati*
University of Connecticut
Dept. of Electrical and Computer Engineering
371 Fairfield Way, U-4157
Storrs, CT 06269-2157
FAX: 860-486-5585
Phone: 860-486-2890
E-mail: krishna@engr.uconn.edu

*To whom correspondence should be addressed: krishna@engr.uconn.edu

Evaluating the Value of Information in the Presence of High Uncertainty*

David Sidoti[1], Diego Fernando Martínez Ayala[1], Xu Han[1], Manisha Mishra[1], Sravanth Sankavaram[1], Woosun An[1], Krishna R. Pattipati[1], and David L. Kleinman[2]

[1]Univ. of Connecticut, Dept. of Electrical and Computer Engineering, Storrs, CT 06269

[2]Naval Postgraduate School, Dept. of Information Science, Monterey, CA 93943

ABSTRACT

Context-driven decision making relies upon a multitude of information, but placing a tangible value on these pieces of information is an upcoming problem in the Information Age. In an application to combating piracy, the US Navy relies upon a probability surface which integrates intelligence, commercial shipping routes, and meteorological and oceanographic (METOC) information, and forms the basis for predicting geo-temporal patterns of pirate presence and attack, commonly referred to as Piracy Attack Risk Surface (PARS). Degradation in the quality of this information lowers the quality of PARS, which forms the basis for counter-piracy surveillance and interdiction asset allocations to geographic regions. In this paper, we investigate the value of PARS in the presence of uncertainty in intelligence and weather when it is used in an asset allocation algorithm that seeks to minimize the probability of success of a pirate attack; the algorithm allocates both interdiction and surveillance assets to deter pirate activities. We perform a sensitivity analysis using hypothetical counter-piracy scenarios to quantify the value of information.

Keywords: Value of information, Counter-piracy, Decision making, Piracy Attack Risk Surface (PARS)

I. INTRODUCTION

Motivation

The Navy has identified a need to quantify the value of data to be delivered to decision makers (DMs) with the aim of eliminating superfluous information, while simultaneously aiding in understanding the context of the pending action [29]. If the value of data were able to be easily measured, distinction of what should and should not be delivered to the DM would be trivial. Emphasis ought to be placed on delivering high-value information to enable faster and more precise decision-making tailored to the situational context. Arranging an information display to achieve higher clarity of information of value to the DMs has been found to lead to faster decisions, thereby lending support to the idea that proper information valuation is a key step towards the successful conduct of warfare [42]. The layered decision architecture proposed in [29], and reproduced in Figure 1, consists of transforming Data $\rightarrow$ Information $\rightarrow$ Knowledge $\rightarrow$ Understanding $\rightarrow$ Wisdom. Sensors and data logging are components of the Data layer, while the concept of value of information in this hierarchy is subsumed in the Knowledge and Information layers. If information with high expected value is identified, it then passes up through the Understanding and subsequent Wisdom layers,

* This work was supported by the U.S. Office of Naval Research under Grant # N00014-12-1-0238.

culminating in DMs' enhanced decision-making. Filtering high-valued information is thus imperative in this architecture in order for the DMs to attain situational awareness and take context-specific courses of action (COA) to achieve superior mission performance.

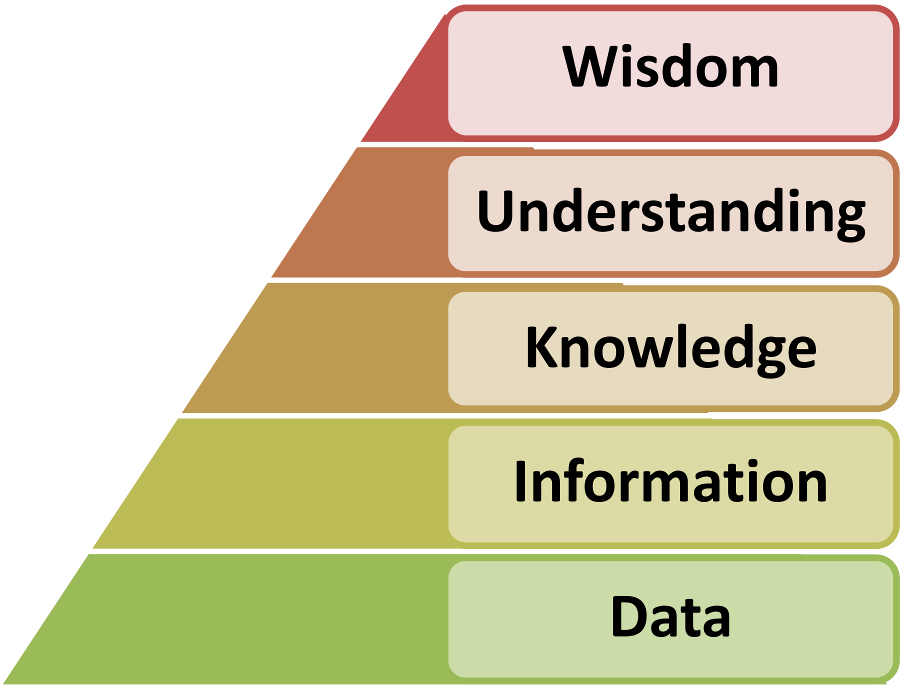

Figure 1: Data to decision making architecture

This paper is concerned with the value of information in the C4ISR domain, specifically counter-piracy in the Gulf of Aden and off the coast of Somalia. In this environment, the Navy uses a probability surface, called Pirate Attack Risk Surface, or PARS. PARS is a grid of cell values that represents the probability of pirate presence (for surveillance) or pirate attack (for interdiction); this surface integrates weather factors, pirate behavior (via intelligence), and commercial shipping routes. PARS map serves as the blueprint for surveillance and interdiction operations for counter-piracy missions. Utilizing the PARS, a dynamic asset allocation to combat the threat of piracy was considered in [2], encompassing both interdiction and surveillance operations.

Given the information gathered and consolidated into the PARS map, the asset allocation problem considered in [2] requires the DMs to choose from a set of available surveillance and interdiction assets to detect and interdict pirates. The DMs' decisions involve which asset to allocate where in the area of responsibility for the next time epoch in order to maximize the probability of detection (for surveillance asset allocation decisions) and minimize the probability of a successful pirate attack (for interdiction asset allocation decisions).

The overall problem, as noted in [7], is that information and data pertaining to the operational environment, mission, and tasks are readily available, but the flow of high-valued information to the DMs requires attention such that the information accessed will provide the most efficient venue for DMs to understand the context of their pending decisions. In this paper, we evaluate the value of PARS itself in the presence of both high and low uncertainty pertaining to weather and intelligence. Quantifying the value of PARS contributes to the overall goal of delivering information of high value to the DMs in a timely manner. Correct evaluation of information in this context is ultimately reflected in the DMs' quick and informed choice of the best COA to identify pirate vessels and prevent/impede attacks on merchant vessels. We also investigate the importance of knowing the COA of each DM as we examine the importance of coordination among assets.

Related Research

The concept of value of information has been applied in a wide array of fields ranging from psychology to medical and machinery diagnosis to dynamic state estimation to oil-drilling [31, 15, 41]. Much of the literature on estimating the value of information stemmed from Bayesian approach to optimal experimental design [31]. In psychology, the well-known Planet Vuma problem [31] hinges on gathering information deemed the "most useful" in order to distinguish between two species. The more useful the information is calculated to be, the higher is its value. Based on the idea of expected utility of asking a question [36], Nelson [31] discusses different approaches for calculating the value of information in the context of the Planet Vuma problem. These include probability gain, information gain, Kullback-Leibler divergence, impact, Bayesian diagnosticity, and log diagnosticity [4, 12-14, 21, 23, 24, 30, 32, 45]. Within a Bayesian decision theoretic framework, Howard [18] formalizes the expected value of a sample of information using the concept of pre-posterior analysis as the expected increase in utility by gaining access to a sample of additional observations.

In the realm of medical and machinery diagnosis, value of information serves as a basis for selecting tests with the maximum information gain per unit cost of the test [44, 33-35, 47]. Performing different tests results in knowledge of different aspects of the patient's (or machine's) condition and so it becomes a primary concern when attempting to determine a diagnosis at minimum cost. As [22] points out, if the cost of testing is greater than the expected value of information, it may not be beneficial for further tests to be administered. Also within the medical field, Dale [8] extends cost-sensitive classification to functional magnetic resonance imaging (fMRI) experiments where he seeks to optimize statistical efficiency by comparing different rates of sampling.

The value of information is used in dynamic sensor scheduling, which has been widely studied in the area of target tracking (e.g., [47]). For linear Gaussian state space systems, one can obtain an analytic solution for the posterior distribution of the system state given the sensor measurements and a scheduling sequence via a Kalman filter [47]. Shakeri *et al*. [47] formulated the sensor scheduling problem subject to a fixed total budget and the cost of individual sensor varying inversely with its measurement variance. They obtained an optimal measurement schedule that minimizes the trace of a weighted sum of the estimation error covariance matrices of a discrete-time vector stochastic process, when the auto-correlation matrix of the process is given. Sub-optimal approaches, based on information-theoretic criteria, have been developed for overcoming the computational intractability of determining the optimal sensor schedule. In the context of sensor networks, Zhao *et al*. [48] formulated the target tracking problem as a sequential Bayesian estimation problem, where the participants for sensor collaboration are determined by minimizing an objective function comprising information utility, e.g., measured in terms of entropy, Mahalanobis distance and the sensor usage cost. In [1], An *et al* employed the auction algorithm for the assignment problem to dynamically allocate sensors to tasks to maximize the cumulative information gain. Similar entropy-based information value metrics are pursued in [2, 9, 17, 20, 25, 37, 38, 46].

The concept of value of information in the context of oil exploration seeks to trade-off the cost of exploration and the value of information. This results in a DM weighing the cost of drilling, rig rates, etc. alongside risks and uncertainty associated with individual oil or gas fields in order to determine if exploration will result in a profit [41]. Milgrom *et al*. [27] examine the negative effects of having such valuable information in a sealed-bid auction for tracts of land sold for oil and gas.

In more recent years, the value of information has found its way into the context of spatial decision making and directed graph models [6, 22]. Bhattacharjya *et al*. [6] present models that compute the value of information when a DM has multiple COAs within the decision space. Krause *et al*. [22] approach the directed graph case by exploiting the use of local rewards and summing for the total reward. Reward functions are presented in the form of residual and joint entropies, enabling a decision-theoretic quantification of the value of information.

To the best of our knowledge, literature pertaining to the value of information is generally lacking in the operations research domain. In [11], the value of information is formalized within a stochastic framework, while [19] concerns itself with *when* information becomes available based upon a DM's decision. Thornley *et al.* [43] consider mission-specific value of information by finding a quantifiable, objective performance metric for commanders to feel confident in a COA.

Value of information, specifically in the context of operations research, is of crucial importance due to ever more restrictive budget constraints within the Navy and the goal of utility maximization. Piracy, especially off the coast of Somalia, has garnered attention within the last few years due to the increase in attacks [2]. In [2] the counter-piracy problem is formulated and solved as a stochastic control problem. In this paper, we find a measure for information value, similar to [43], via the PARS. By solving the stochastic control problem, we use a quantifiable metric known as probability of interdiction (also referred to as interdiction gain) to justify the effectiveness of the PARS and to contrast our solutions in the presence of high uncertainty. In the context of C4ISR, such a performance metric is available through the algorithm in [2] when the probability of interdiction is used as a reward metric for the asset. We use cumulative probability of interdiction over all assets as a metric for quantifying the value of information in the context of counter-piracy.

Organization of the Paper

This paper is organized as follows. In Section II, we introduce the components and significance of the PARS and propose a probability of interdiction metric to quantify the value of PARS. We also consider the impact of asset coordination on the interdiction metric. In Section III, the value of PARS is evaluated using a hypothetical piracy scenario and the results of sensitivity analysis are discussed. Section IV concludes with our summary and future research directions.

II. PROBLEM FORMULATION AND APPROACH

Pirate Attack Risk Surface (PARS)

The Pirate Attack Risk Surface is a set of heat maps that are output by the decision making software for naval operations discussed in [2] and is calculated using multiple pieces of information. A rectangular ocean area is discretized into grid of cells, wherein each cell consists of a value indicating the probability of a pirate attack within the cell's geographical location. Each cell's probability is calculated based on intelligence regarding pirate behavior, shipping routes, and METOC conditions [16]. This probability of attack results from the multiplication of three components: 1) the probability that pirates will be in the area of interest (AOI), 2) a probability field apropos of vulnerability of shipping activity in the AOI, and 3) the probability that a pirate attack is possible given any possible environmental impacts. A contour map output of PARS is shown in Figure 2.

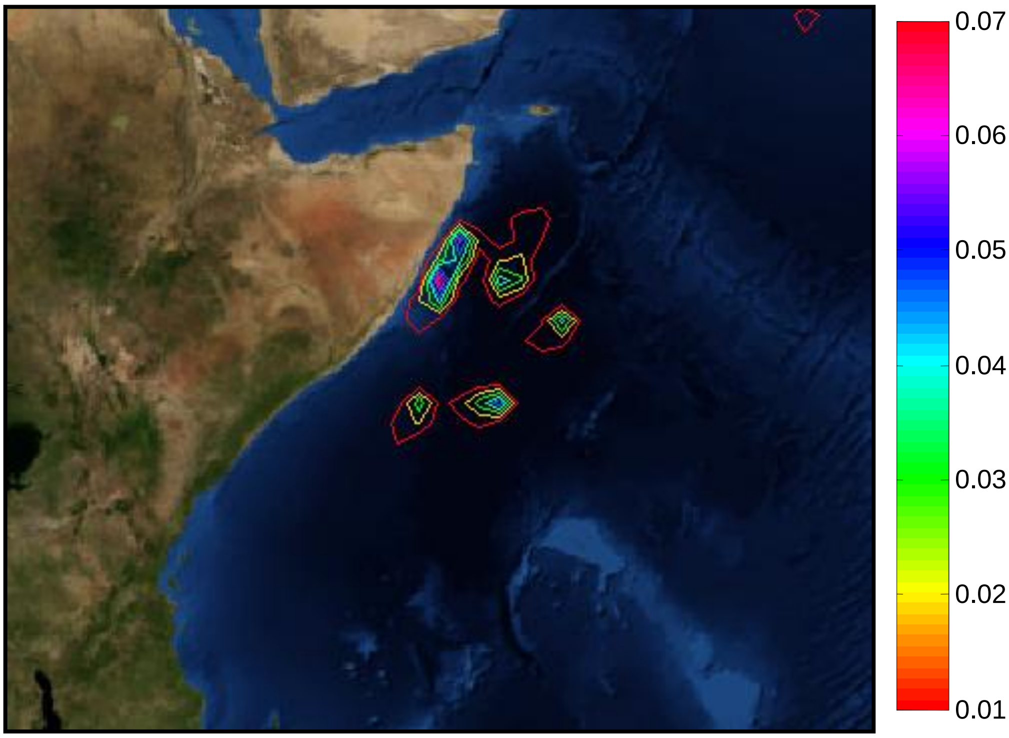

Figure 2 Example of a contour map of a PARS output where the AOI is off the coast of Somalia

These surfaces serve as a visual aid for probable sites of attack and are used as input for the counter-piracy stochastic control model in [2]. The planning process is repeated on a typically 12 to 24 hour cycle, wherein the aforementioned three components are subsequently updated resulting in a new PARS. The counter-piracy model operates on a rolling horizon planning assumption (e.g., 72 hours); thus the set of PARS maps are generated over a finite horizon, each map corresponding to a different time interval, and then regenerated upon entering the next time interval.

Probability of Interdiction/ Interdiction Gain

To derive the measure of performance used, interdiction gain, we first briefly discuss the interdiction formulation in [2] and then explain the metric in detail. We define a time period k, to be the beginning of a 12 hour duration between updates to the PARS and assume the current time period to be denoted by time $k = 0$. At $k = 0$, the DMs decide on an asset allocation for the next K periods, $k = 1, 2, \ldots, K$. The DMs decide on a course of action based on the set of forecasted PARS maps. For our purposes, we are concerned only with interdiction assets, $x_i(k)$, where $i \in I_k$, where $x_i(k)$ denotes the location of an interdiction asset indexed by i at time epoch k, and I_k denotes the set of available interdiction assets at time epoch k. Interdiction assets are able to be assigned during time periods $k > 0$. Note that k is a relative time index.

At each time epoch k, each available interdiction asset $x_i(k)$ has a probability of interdicting a pirate in cell $g \in G$, G denoting the set of PARS grid cells in the AOI. We solve for each 12 hour cycle and use the previous optimization's results as initial conditions for the succeeding time epoch. Interdiction assets (typically surface ships) with speed v_i, are assumed to have a helicopter onboard where the helicopter's speed, v_i^h and its launch delay time, t_i^h are considered in the interdiction probability, denoted $PI_i(x_i(k), g)$. If we suppose the time to interdict a piracy event to be τ, the probability of interdiction is given by,

$$PI_i\left(x_i\left(k\right),g\right)=\begin{cases}\dfrac{2r\left(i,\tau\right)}{dist\left(x_i\left(k\right),g\right)}, & r\left(i,\tau\right)<\dfrac{dist\left(x_i\left(k\right),g\right)}{2}\\ 1, & r\left(i,\tau\right)\geq\dfrac{dist\left(x_i\left(k\right),g\right)}{2}\end{cases}, \tag{1}$$

where $dist(x_i(k), g)$ is the Euclidian distance from cell g to the location of asset $x_i(k)$. Here, $r(i,\tau)$ is the distance that will be covered during τ and takes on the following values conditioned upon each helicopter's required time to launch.

$$r\left(i,\tau\right)=\begin{cases}v_i\tau, & \tau\leq t_i^h\\ v_i t_i^h+v_i^h\left(\tau-t_i^h\right), & \tau>t_i^h\end{cases}, \tag{2}$$

We then are able to use the cumulative interdiction probability,

$$CPI=\sum_{k=0}^{K}\sum_{i\in I_k}\sum_{g\in G}PI_i\left(x_i\left(k\right),g\right) \tag{3}$$

as a benchmark to evaluate the value of PARS. Our approach allows us to compare the probability of interdiction accumulated over the time horizon in both high and low uncertainty cases. The difference between the cumulative interdiction gains of two qualitatively different PARS maps allows us to evaluate the importance of the quality of PARS maps in counter-piracy operations.

Asset Coordination

In an ideal environment, each decision is known by each asset globally; information is easy to acquire and so DMs can choose a COA based on a COA already chosen by another DM. In [2], no assets are allowed to enter the same cell in the same time epoch and so each path chosen is known by all of the DMs. Here, we define a.) *coordinated assets*: each asset is aware of the paths of the remaining assets, b.) *partially uncoordinated assets*: some of the assets are aware of the paths of the remaining assets, c.) *uncoordinated assets*: none of the assets are aware of the paths of the remaining assets. When fully coordinated, the environment can be explored efficiently with no overlap between interdiction vessels. We compare the scenarios of coordinated, partially uncoordinated, versus uncoordinated assets using the cumulative interdiction probability as a metric for comparison. We solve the counter-piracy problem for one asset, and subsequently for the remaining assets to simulate an asset uninformed about other assets' routes. Doing this for each asset simulates a scenario where there exists a complete breakdown in communication amongst DMs. Evaluating the cumulative probability of interdiction in each instance quantifies the importance of coordination amongst DMs.

III. COMPUTATIONAL RESULTS

Mission Scenario and Results

We use the same area of interest (AOI) as in [2], where the region is discretized into a grid of cells corresponding to available METOC forecasts. The cells are squares of 0.8-arcdegree-side length and merge to become a 43×51 cell grid. Here, we consider the placement of multiple interdiction vessels of identical capability over a finite planning horizon of three days. The asset capabilities input to the model are similar to those of a Ticonderoga-class guided missile cruiser. See Table I for the asset characteristics. Note in Table I, "H" under "Asset Types Carried" signifies a helicopter (or multiple helicopters) on board.

Table I: Interdiction vessel characteristics

Max Speed (km/h)	Range (km)	Home Base	Asset Types Carried	# of Assets Carried	Surface Radar Range (km)
60	6100	US	H	2	92

We solve the interdiction problem for two cases: 1) PARS in the presence of high uncertainty, and 2) PARS in the presence of low uncertainty; we conduct a sensitivity analysis alongside our results and find the routes and the corresponding interdiction gains for four scenarios of available interdiction assets: 2, 4, 7, and 10 vessels.

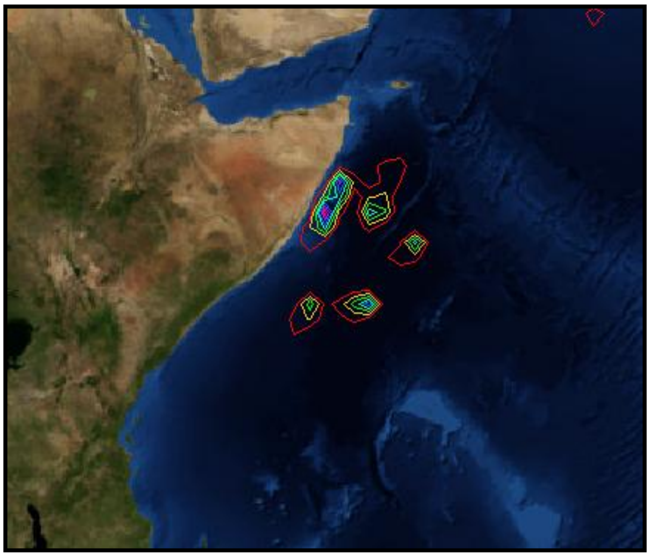

(a) Low uncertainty

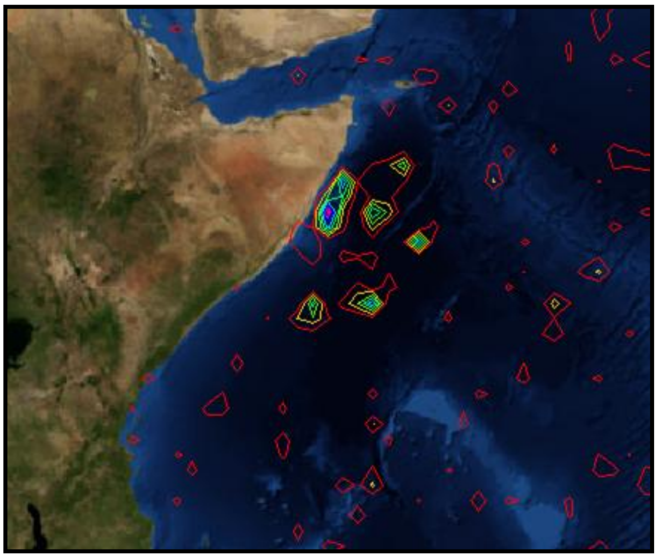

(b) Medium uncertainty

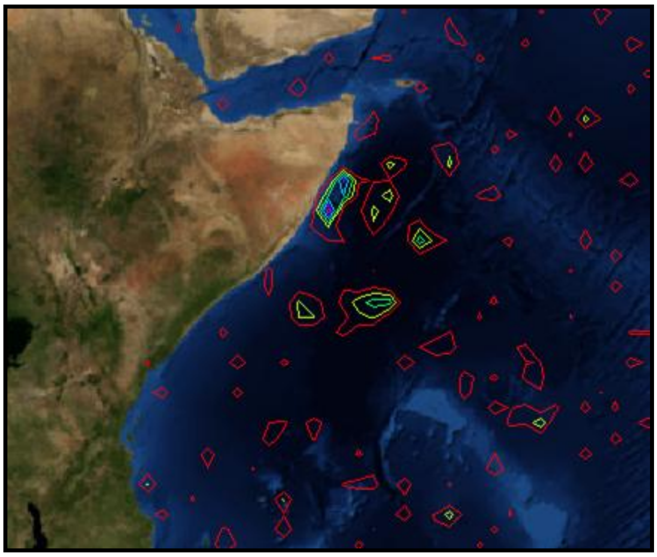

(c) High uncertainty

Figure 3: Degraded PARS by adding Gaussian noise to the Best Case PARS

To capture the uncertainty aspect in PARS, we add zero-mean Gaussian noise to each individual cell in the PARS, and increase the variance as the forecasting time is further into the future. We compare two surfaces for each quantity of assets – one using the original PARS, and the other using a PARS with noise with a variance of $(0.04)^2$ added at the current time epoch wherein the variance increases by a factor of two for each subsequent forecasting time epoch (e.g. $(0.04)^2$, $(0.08)^2$, $(0.16)^2$, etc). This noise is added to the PARS via a moving window of length 6 – two 12-hour updates each day for 3 days. Due to the increasing variance, the signal-to-noise ratio decreases over the time horizon as shown in Table II.

Table I: SNR (dB) for the PARS with added uncertainty

k	**1**	**2**	**3**	**4**	**5**	**6**
SNR	16.2872	13.9694	13.5497	8.2313	8.6664	4.9296

Examples of degradation in the quality of PARS maps are shown in Figure 3. Assets are positioned uniformly such that there is no obtainable interdiction gain on day 0 before the first update in the PARS.

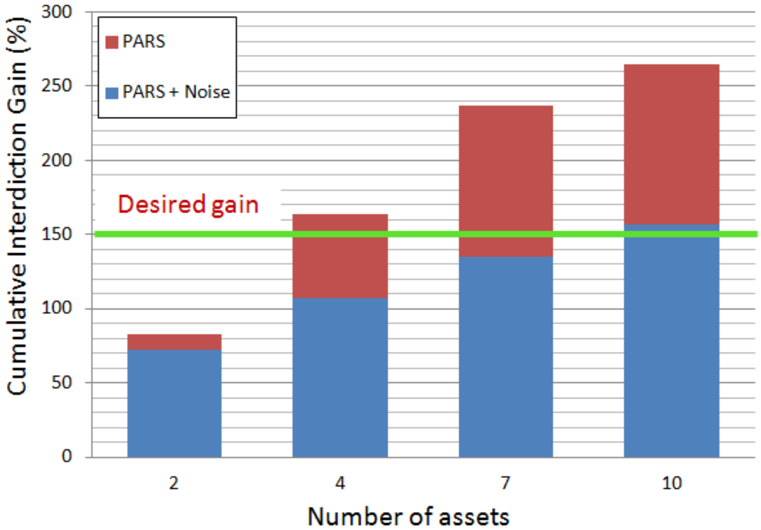

Figure 4: Cumulative interdiction gain (%) for PARS

Table II: Comparison of Interdiction Gains

# of Assets	**Cumulative Probability of Interdiction (%)**	
	PARS + Noise	**PARS**
2	72.319	82.730
4	107.45	163.95
7	134.87	236.54
10	156.63	264.46

Figure 4 shows the cumulative interdiction gain as a function of the number of assets comparing the PARS with low uncertainty with that of high uncertainty (noise with a variance that increases by a factor of two for each subsequent time epoch, as mentioned earlier). Summing the PARS, there are three pirates in the AOI in this hypothetical scenario, thus a 300% interdiction gain corresponds to successfully interdicting all three pirate attacks. The sensitivity analysis results for each case are shown in Table III. Results are obtained by averaging the results over 100 Monte Carlo runs each with the exception of when 10 interdiction assets are available which, due to computational complexity, was averaged over 10 Monte Carlo runs only.

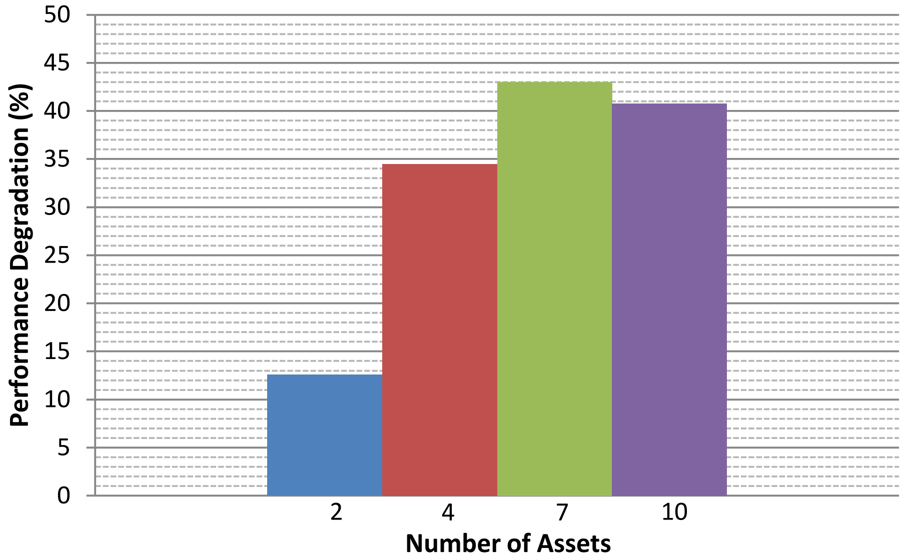

Figure 5: Performance degradation (%) due to PARS maps in the presence of high uncertainty

Value of Information Analysis

In order to quantify the value of PARS, we set a desired interdiction gain of 150% (equating to an overall 50% successful interdiction rate) over the course of 6 time epochs, or 3 days, for the sake of simplicity and ease of visual comparison. As shown in Figure 4, the interdiction asset allocation plan with the original PARS required only *four assets*, where as a plan with highly uncertain PARS required *ten assets*; this translates to a 150% increase in asset requirements for the same desired interdiction gain. The difference between having a PARS with low uncertainty and a PARS with high uncertainty can mean using up to *six assets less*. Interpreting our results a step further, this means that having a "good" PARS can be extremely economical, ultimately allowing the DMs to forego the operating costs of six vessels. Figure 5 shows that these results are due to the rate of increase in the cumulative interdiction gain for PARS with high uncertainty with respect to the number of assets being less than that of PARS with low uncertainty. Performance degradation was calculated via

$$\frac{CPI_H - CPI_L}{CPI_L} \times 100\%, \tag{4}$$

where CPI_H denotes the cumulative probability of interdiction obtained when the PARS has high uncertainty and CPI_L denotes the contrary case. Thus, for two assets, using PARS with high uncertainty results in approximately 12% degradation in performance. Using four assets, this gap widens to 34%. We see that performance is most disparate when seven assets are used (43%), signifying performance degradation as a consequence of adding high uncertainty to the forecast. Tables IV and V show the cumulative interdiction gain obtained for each time epoch using both the degraded and the best case PARS, respectively.

Table III: Interdiction Gains using the Degraded PARS

Time *k* / # of Assets	1	2	3	4	5	6
2	0	0.0284	0.1222	0.2680	0.4817	0.7231
4	0	0.0321	0.1570	0.3477	0.6536	1.0744
7	0	0.0369	0.1609	0.3973	0.7689	1.3486
10	0	0.0557	0.2118	0.5058	0.9446	1.5663

Table IV: Interdiction Gains using the Best Case PARS

Time *k* / # of Assets	1	2	3	4	5	6
2	0	0.0142	0.1539	0.3750	0.5746	0.8273
4	0	0.0259	0.3038	0.7680	1.1855	1.6395
7	0	0.0742	0.4680	1.0418	1.6841	2.3654
10	0	0.0742	0.5311	1.1493	1.8511	2.6446

Value of Coordination

Using our findings from our value of information analysis, we examine coordination amongst four assets using the best case PARS to investigate the cumulative interdiction gain as a function of asset coordination. Utilizing the same locations as before when four assets were allocated, we found full asset coordination to be necessary in order to obtain a desired cumulative interdiction gain of 150%. As shown in Figure 6, lack of asset coordination lowers cumulative interdiction gain by 10% or more and results in a failure to obtain our desired gain. Coordinating interdiction strategies improve gain by as much as 22% over the uncoordinated case.

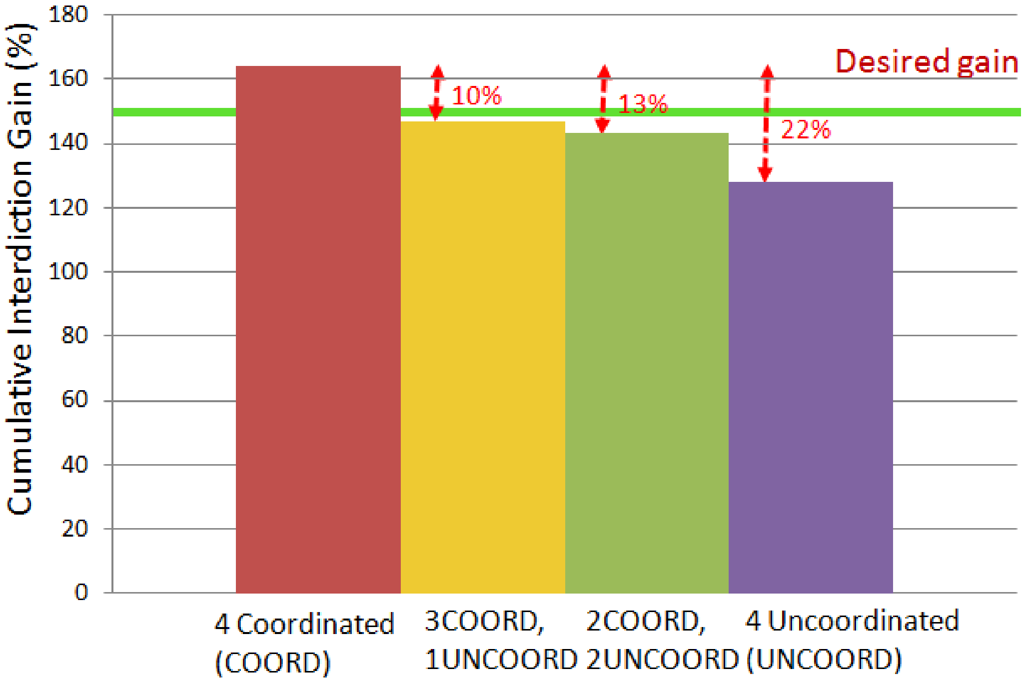

Figure 6: Cumulative interdiction gain versus asset coordination for four assets.

IV. CONCLUSION AND FUTURE WORK

This paper considered the problem of quantifying mission-specific value of information using counter-piracy operations as mission context. The interdiction asset allocation algorithm in [2] was used to quantify the value of PARS maps by considering various levels of uncertainty in these maps ranging from low to high uncertainty. The experimental results demonstrate the importance of PARS and the degradation in solution quality in the presence of high uncertainty. An analysis of coordinated versus uncoordinated assets was conducted, where a fully coordinated asset scenario satisfied the desired gain requirement; this is commensurate with our intuition.

Future research will shift the focus of our model and analysis to the counter-smuggling missions in the East Pacific and Caribbean oceans. In this environment, other variations of mission performance metrics are being developed to quantify the value of PARS maps. These include metrics such as the expected amount of drugs interdicted and expected number of vessels interdicted. In addition, we will explore other measures of information value based on Bayesian diagnosticity, impact, information gain, and other Bayesian OED framework theories discussed in [31].

REFERENCES

[1] W. An, S. Singh, K. R. Pattipati, S. Gokhale, and D. Kleinman, "Dynamic Scheduling of Multiple Hidden Markov Model-based Sensors", *Journal of Advances in Information Fusion*, Vol. 3, No. 1, pp. 33-49, 2008.

[2] W. An, D. Ayala, D. Sidoti, M. Mishra, X. Han, and K. R. Pattipati, "Dynamic asset allocation approaches for counter-piracy operations," in *Information Fusion (FUSION), 2012 15th International Conference on*, pp. 1284–1291, 2012.

[3] P.-E. Back, L. Rosén, and T. Norberg, "Value of information analysis in remedial

investigations," *Ambio*, vol. 36, no. 6, pp. 486–93, 2007.

[4] J. Baron, *Rationality and intelligence*. Cambridge, England: Cambridge University Press, 1985.

[5] T. Bayes, "An essay towards solving a problem in the doctrine of chances," *Philosophical Transactions of the Royal Society of London*, vol. 53, pp. 370–418, 1763.

[6] D. Bhattacharjya, J. Eidsvik, and T. Mukerji, "The Value of Information in Spatial Decision Making," *Mathematical Geosciences*, vol. 42, no. 2, pp. 141–163, 2009.

[7] K. Byers, "Situational Awareness for Surveillance and Interdiction Operations (SASIO): Tactical Installation Protection," M.S. thesis, Dept. of Operations Research, Naval Postgraduate School, Monterey, CA, 2010.

[8] A. Dale, "Optimal experimental design for event-related fMRI.," *Human brain mapping*, vol. 8, no. 2–3, pp. 109–14, 1999.

[9] P. Dodin, J. Verliac, V. Nimier, "Analysis of the multisensor multitarget tracking resource allocation problem," in: *Proceedings of the International Conference on Information Fusion*, 2000, pp. WeC1-17–22.

[10] L. Esher, S. Hall, and E. Regnier, "Simulating pirate behavior to exploit environmental information," *Proceedings of the 2010 Winter Simulation Conference*, vol. 6, pp. 1330–1335, 2010.

[11] V. Goel and I. E. Grossmann, "A Class of stochastic programs with decision dependent uncertainty," *Mathematical Programming*, vol. 108, no. 2–3, pp. 355–394, 2006.

[12] I. J. Good, *Probability and the weighing of evidence*. New York: Griffin, 1950.

[13] I. J. Good, "Explicativity, corroboration, and the relative odds of hypotheses." *Synthese*, vol. 30, pp. 39–73, 1975.

[14] I. J. Good, *Good thinking*. Minneapolis: University of Minnesota, 1983.

[15] I. J. Good and W. I. Card, "The diagnostic process with special reference to errors.," *Methods of information in medicine*, vol. 10, no. 3, pp. 176–88, 1971.

[16] J. A. Hansen, D. Hodyss, C. H. Bishop, and W. Campbell, "Coupled INTEL/METOC Risk Assessment," Patent Publication # 20120095946, April 19, 2012.

[17] K. J. Hintz, E.S. McVey, "Multi-process constrained estimation," *IEEE Transactions on Systems, Man, and Cybernetics,* vol. 21, no. 1, pp. 434–442, 1991.

[18] R. Howard, "Information Value Theory," *IEEE Transactions on Systems Science and Cybernetics*, vol. 2, no. 1, pp. 22–26, 1966.

[19] T. W. Jonsbråten, R. J. B. Wets, and D. L. Woodruff, "A Class of Stochastic Programs with Decision Dependent Random Elements," *Annals of Operations Research*, vol. 82, pp. 83–106, 1998.

[20] K. Kastella, "Discrimination gain to optimize detection and classification," *IEEE Transactions on Systems, Man, and Cybernetics, Part A: Systems and Humans,* vol. 27, no. 1, pp. 112–116, 1997.

[21] J. Klayman and Y.-W. Ha, "Confirmation, disconfirmation, and information," *Psychological Review*, vol. 94, pp. 211–228, 1987.

[22] A. Krause and C. Guestrin, "Optimal value of information in graphical models,"

Journal of Artificial Intelligence Research, vol. 35, pp. 557–591, 2009.

[23] S. Kullback and R.A. Liebler, “Information and sufficiency,” *Annals of Mathematical Statistics*, vol. 22, pp. 79–86, 1951.

[24] D.V. Lindley, “On a measure of the information provided by an experiment,” *Annals of Mathematical Statistics*, vol. 27, pp. 986–1005, 1956.

[25] G. A. McIntyre, K.J. Hintz, “An information theoretic approach to sensor scheduling,” *Proceedings of the SPIE International Symposium on Aerospace/Defense Sensing and Control*, Orlando, FL, vol. 2755, pp. 304–312, 1996.

[26] L. Meier, J. Perschon and R. M. Dressier, “Optimal control of measurement subsystems,” *IEEE Transactions on Automatic Control*, vol. 12, pp. 528-536, 1967.

[27] P. Milgrom and R. J. Weber, “The value of information in a sealed-bid auction,” *Journal of Mathematical Economics*, vol. 10, no. 1, pp. 105–114, 1982.

[28] R. Mirshak, “Ship Response Capability Models for Counter-Piracy Patrols in the Gulf of Aden,” *DRDC CORA TM* 2011-139, 2011.

[29] J. G. Morrison, (2011, December 1) *Data to Decisions or Decisions to Data? A Human System Perspective on D2D*, [Online]. Available: http://www.ictas.vt.edu/cnavs/presentations/data_to_decision/4morrison.pdf.

[30] J. Nelson, “Finding useful questions: on Bayesian diagnosticity, probability, impact, and information gain.,” *Psychological review*, vol. 112, no. 4, pp. 979–99, 2005.

[31] J. Nelson, “Towards a rational theory of human information acquisition,” *The probabilistic mind: Prospects for rational models of cognition*, pp. 143–163, 2008.

[32] R. S. Nickerson, “Hempel’s paradox and Wason’s selection task: Logical and psychological puzzles of confirmation,” *Thinking and Reasoning*, vol. 2, pp. 1–32, 1996.

[33] K. R. Pattipati and M.G. Alexandridis, “A Heuristic Search and Information Theory Approach to Sequential Fault Diagnosis,” *IEEE Transactions on Systems, Man, and Cybernetics*, pp. 872-887, 1990.

[34] V. Raghavan, M. Shakeri and K.R. Pattipati, “Optimal and Near-optimal Test Sequencing Algorithms with Realistic Test Models,” *IEEE Transactions on Systems, Man, and Cybernetics: Part A - Systems and Humans*, vol. 29, no. 1, pp. 11-27, 1999,.

[35] V. Raghavan, M. Shakeri and K.R. Pattipati, “Test Sequencing Problems Arising in Test Planning and Design for Testability,” *IEEE Transactions on Systems, Man, and Cybernetics: Part A - Systems and Humans*, vol. 29, no. 2, pp. 151-163, 1999.

[36] L. J. Savage, *The Foundations of Statistics*. New York: Wiley Subscription Services, Inc., A Wiley Company, 1954.

[37] W. Schmaedeke, “Information based sensor management, in: Signal Processing, Sensor Fusion, and Target Recognition II,” Orlando, FL, *Proceedings of the SPIE 1955*, 156–164, 1993.

[38] W. Schmaedeke, K. Kastella, “Information based sensor management and IMMKF,” *Proceedings of the SPIE International Conference on Signal and Data Processing of Small Targets*, Orlando, FL, pp. 390–401, 1998.

[39] M. Shakeri, K. R. Pattipati and D. L. Kleinman, “Optimal measurement scheduling for

state estimation," *IEEE Transactions on Aerospace and Electronic System*, vol. 31, pp. 716-729, 1995.

[40] L. A. Slootmaker, "Countering Piracy with the Next-Generation Piracy Performance Surface Model," DTIC Document, 2011.

[41] P. C. Smalley, S. H. Begg, M. Naylor, S. Johnsen, and A. Godi, "Handling risk and uncertainty in petroleum exploration and asset management: An overview," *AAPG Bulletin*, vol. 92, no. 10, pp. 1251–1261, 2008.

[42] M. St. John, H. S. Smallman, D. I. Manes, B. A. Feher, and J. G. Morrison, "Heuristic automation for decluttering tactical displays.," *Human Factors*, vol. 47, no. 3, pp. 509–25, 2005.

[43] D. J. Thornley, R. J. Young, and J. P. Richardson, "Toward mission-specific service utility estimation using analytic stochastic process models," in *Proc. SPIE 7352, Intelligent Sensing, Situation Management, Impact Assessment, and Cyber-Sensing*, 2009, vol. 7352, pp. 1-9, 2009.

[44] P. Turney, "Cost-sensitive classification: Empirical evaluation of a hybrid genetic decision tree induction algorithm," *Journal of Artificial Intelligence Research (JAIR)*, vol. *2*, pp. 369–409, 1995.

[45] G. L. Wells and R. C. L. Lindsay. "On estimating the diagnosticity of eyewitness nonidentifications," *Psychological Bulletin*, vol. 88, pp. 776–784, 1980.

[46] N. Xiong, P. Svensson, "Multi-sensor management for information fusion: issues and approaches," *Information Fusion*, vol. 3, no. 2, pp. 163-186, 2002.

[47] H. Ying and K. P. Chong, "Sensor scheduling for target tracking in sensor networks," *IEEE Conference on Decision and Control*, vol. 1, pp. 743-748, 2004.

[48] F. Zhao, J. Shin and J. Reich, "Information-driven dynamic sensor collaboration," *IEEE Signal Processing Magazine*, vol. 19, pp. 61-72, 2002.