18th International Conference on Information Fusion
Washington, DC - July 6-9, 2015

Dynamic Resource Management and Information Integration for Proactive Decision Support and Planning

Manisha Mishra, David Sidoti, Diego Fernando Martinez Ayala, Xu Han, Gopi Vinod Avvari, Lingyi Zhang, and Prof. Krishna R. Pattipati,
Department of Electrical and Computer Engineering
University of Connecticut, Storrs, CT, U.S.A.
Email: krishna@engr.uconn.edu

Dr. Woosun An[1], Dr. James A. Hansen[2], and Prof. David L. Kleinman[3]
[1]SIMNET Co. Ltd., South Korea
[2]Naval Research Lab, Monterey, CA
[3]Naval Postgraduate School, Monterey, CA
Email: dlkleinm@nps.edu

***Abstract*—Major challenges anticipated by future mission planners comprise automated processing, interpretation, and development of intelligent decisions using large volumes of dynamically evolving structured and unstructured data, while simultaneously decreasing the time necessary to plan and re-plan. Motivated by the need to seamlessly integrate automated information processing and resource management for proactive decision-making and execution in a highly adaptive network-centric environment, we propose a) surveillance and interdiction algorithms for dynamic resource management; b) distributed and collaborative mixed-initiative multi-level resource allocation algorithms to allocate hierarchically-organized assets to process inter-dependent tasks and goals; and c) quantifying the value of information in order to accomplish mission objectives. The decision support concepts and algorithms discussed in this paper seek to maximize the efficiency of information transactions in mission planning/re-planning processes by achieving shared situational awareness and increased mission effectiveness. We specifically focus on the dynamic decision making processes associated with planning in a broad range of maritime operations.**

***Index Terms*—Resource management, surveillance, interdiction, value of information, multi-level resource allocation**

I. Introduction

Future mission planning environments are expected to be complex and distributed due to rapid advances in cyber-physical systems and the ubiquitous use of intelligent heterogeneous assets (e.g., unmanned aerial and undersea systems) to operate from ships/sea bases. These assets (resources), in contrast to manned systems, offer a unique set of capabilities, including ultra-long endurance and high-risk mission acceptance. Additionally, these heterogeneous resources provide the ability to automatically collect massive amounts of structured, unstructured, and semi-structured data, which form the information base for decision making. Currently, the existing decision support tools are inundated with too much data and not with enough *information*; this increases cognitive demands on the decision makers (DMs) by diverting their attention to irrelevant data and by driving them to perform competing and conflicting mission tasks without ascertaining the current mission context. Channelized attention of DMs sometimes leads to operational mishaps – as an example, operators flying a Predator were so focused on completing the assigned search and rescue task that they failed to notice that an unmanned aircraft was headed toward a mountain; the aircraft was destroyed on impact and damage was estimated to be $3.9M. Moreover, eight of the soldiers who were to be provided air support by the Predator were killed. If information alerts on situational awareness and tasking had been appropriately allocated to the right operators on the team, those lives may have been saved [19]. Therefore, it becomes imperative to systematically analyze, process and interpret the collected data for presenting relevant information to DMs in a timely manner and making effective decisions even under dynamic, uncertain, and challenging mission conditions (e.g., changes in mission goals, environment, assets and tasks).

In addition to information processing, it is also crucial to dynamically allocate scarce and expensive resources to collect decision-relevant information that has the potential to increase the probability of mission success. For example, in counter-smuggling missions, which involve surveillance (to search, detect, track and identify potential drug trafficking vessels) and interdiction operations (to intercept, board, investigate and potentially apprehend suspects), the *joint* problem of dynamically allocating surveillance and interdiction assets to optimally trade-off exploration versus exploitation in thwarting potential smuggling activities under different operational contexts is very challenging [12]. Given the probability of activity of the smugglers, a function of the intelligence and meteorological and oceanographic (METOC) information, that predicts where the smugglers may transit, the key operational decisions to be made are the following. Where should the surveillance assets and the surface interdiction vessels be allocated to maximize the probabilities of detection and interdiction, respectively?; When and how should the assets transit to the search location?; How to dynamically route and coordinate surveillance and interdiction assets under uncertain weather conditions?; How to best allocate resources when there is missing or only

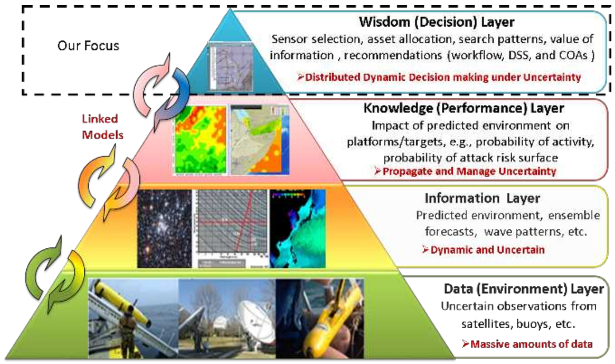

Fig. 1. Battlespace on Demand (BonD) Framework [2]

historical information (also referred to as "flow") regarding the targets is available; How often should re-planning be done and how should this information be conveyed to the DMs unobtrusively? Similar problems arise in counter-piracy, anti-submarine warfare (ASW), search and rescue (SAR), and unmanned aerial vehicle (UAV) missions.

A. Technical Challenges

We posit that these operational challenges can be abstracted into the following key technical challenges, which will be faced by future DMs in planning their missions: 1) The synergy created by a blended force of manned and unmanned assets will be of paramount importance to the joint forces with the concomitant problem of finding solutions for efficient routing, scheduling, dynamic coordination and multi-level mission planning in a contested mission space; 2) With increasing sensing capabilities, identifying, extracting and fusing arrays of information embedded in huge volumes of data and assuring mission success in the face of cyber threats and uncertainty will be computationally expensive and challenging. 3) The uncertain and dynamically changing patterns of potential threats and conflicts in today's world require a strong military capability that is sufficiently proactive to execute a full range of operations – from normal peacetime operations to major combat operations, humanitarian assistance/disaster relief and asymmetric threats such as piracy and terrorism. The term proactive extends beyond adaptability; its virtues range from *responsiveness, congruence with mission, robustness, innovativeness, flexibility,* and *anticipation* to *resiliency*.

Currently, a conceptual framework for mission planning is the Battlespace on Demand (BonD) [2]. This overarching decision support framework allows the commanders to access, assess, and use environmental information to make better decisions, to improve their awareness of their present environmental context, and to improve information sharing among commanders, thereby improving the situational awareness of the battlespace [2]. Additionally, the framework provides a systematic approach to convert knowledge of the forecasted oceanographic environment to be incorporated into war-fighting and shaping decisions. The BonD framework, as shown in Figure 1 includes four layers: Data, Environment, Performance, and Decision. The Data Layer (Tier 0) incorporates massive amounts of structured and semi-structured data from various heterogeneous sensors (e.g., satellites, buoys, etc.). The Environment Layer (Tier I) allows the commanders to assess the present and future states of the oceanographic environment. In the Performance Layer (Tier II), the impact of the current and forecasted environment on individual sensors and weapon platforms, in the form of Performance surfaces, is evaluated. In the Decision Layer (Tier III), the commanders translate the knowledge of the current and forecasted physical environment, along with its uncertainty and its impact on assets (i.e., sensors and weapons), into mission risk and meaningful courses of action. Currently, the Tier III Decision layer primarily requires a manual capability with experienced, highly-trained personnel for superior mission performance. Here, enabling commanders to make informed decisions across a range of maritime operations and skill levels of personnel is a key element to enhance mission effectiveness and planning.

Thus, motivated by the need to seamlessly integrate automated information processing, and dynamic resource management in a highly adaptive network-centric environment, we seek to introduce dynamic decision support concepts and algorithms to maximize the efficiency of information transactions in multi-mission planning/re-planning processes to achieve shared situational awareness and improved mission effectiveness. The proactive mission planning algorithms discussed in this paper dynamically invoke plans as a function of emerging events, while readily adapting to meet unfolding events by monitoring the outcomes of many of its previous decisions, and re-plan, if warranted.

The remainder of the paper is organized as follows. Section II discusses a dynamic asset allocation problem motivated by operational concerns, such as counter-smuggling and counter-piracy operations. In Section III, a distributed and collaborative resource management for Maritime Operations Center planning is discussed. In Section IV, we discuss methods to quantify the value of information for decision making. Finally, the paper concludes with a summary of key findings and future research directions in section V.

II. Dynamic Resource Management

The problem of dynamic resource management under uncertainty, arising in surveillance and interdiction operations, may be viewed as a moving horizon stochastic control problem, as shown in Figure 2. In the context of a counter-piracy/counter-smuggling mission, the problem is to efficiently allocate a set of heterogeneous sensing and interdiction assets to maximize the probability of pirate/smuggler detection and interdiction, subject to mission constraints by integrating information, such as INTEL, weather, asset availability, asset capabilities (e.g., range, speed), sensor management, and asset assignment (e.g., many sensors may need to be coordinated to obtain a better picture of the situation). These problems are NP-hard and can be mapped closely to the decision support layer of the aforementioned BonD framework in the previous section.

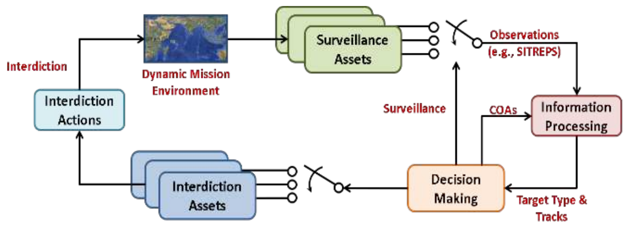

Fig. 2. Dynamic Resource Management Problem as a Stochastic Control Problem [3]

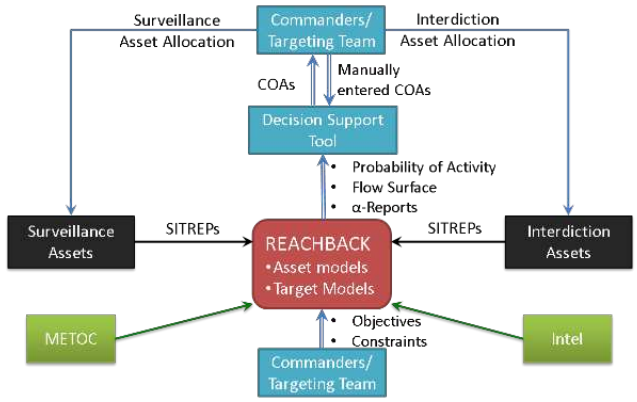

Fig. 3. Information Flow in a Decision Making Process [12] [18]

Probabilistic information on smuggling/piracy activity is generated in the form of color coded heat maps based on INTEL and METOC (e.g., wind speed and direction, wave heights, and ocean currents) information as shown in Figure 3. This information is interpreted in the form of probability of activity (PoA)/pirate attack risk surfaces (PARS) surfaces. The PoA/PARS [3][11] surfaces are represented in the form of longitude-latitude-specific probability mass functions, indicating the behavior/activity of each smuggler in conjunction with a set of descriptors for each smuggler presented in the form of active/pending cases. The PoA/PARS surfaces are akin to the "information state" in the stochastic control setting and thus forms the "sufficient statistics" for the asset allocation and scheduling process. The operational planning problems for counter-smuggling and counter-piracy operations are similar in that they are both characterized by multiple (distinct) targets and assets; the inter-temporal constraints combined with inter-temporal correlations in information, and coupled search and interdiction with different time-scales. In the next section, we discuss the problem formulation of counter-piracy mission which can be easily extended to counter-smuggling operation, as well.

A. Interdiction Operation

We consider a moving horizon planning problem, where each time period $k (k = 1, 2, ..., K)$ is of length Δ, (e. g., 1- 24 hours), denoting the time interval between the PARS updates and the status information on available assets. DMs plan at the current time period $(k = 0)$ as to where the assets are to be positioned for the next K periods, $k = 1, 2, ..., K$ using the PARS surface and asset status as the basis for asset allocation at the beginning of every planning period. A forecast of the PARS surfaces containing the evolution of pirate probabilities (target dynamics) is given to the DM at the beginning of each planning period based on the latest information [3]. In the parlance of stochastic control, PARS constitutes the "sufficient statistics"("information state") for decision making [15]-[17]. Here, we assume that the target dynamics do not change between two consecutive PARS updates i.e. Δ.Thus, as Δ approaches zero the PARS surfaces provide a real time update of the targets.

The set of interdiction assets that are assigned during period $k > 0$ is denoted as I_k, $k = 1, 2, ..., K$. The area of responsibility, G, is partitioned into cells denoted by $g \in G$. PARS, updated at the end of every time period, provides the probability of pirate attack in cell g during time period k, denoted by $PA(g, k)$. The cell location of asset i during time period k is denoted as $x_i(k)$. Decisions that are made at $k = 0$ correspond to the future positioning of the available assets. Thus, at time $k = 0$, the decision variables are:

$$\mathbf{U} = \{x_i\,(k)\ \forall k = 1, 2, .., K,\ \ \forall i \in I_k\} \tag{1}$$

Given $x_i(k)$, each asset can traverse a set of reachable cells $R_i(x_i(k)) \subset G$ in time period $k+1$ depending on METOC effects at time k and the asset speed. Thus, $R_i(x_i(k))$ can be a function of k, but does not show its explicit dependence on k for simplicity of notation. In order to select the optimal policy $\mathbf{U}^*$ over the planning horizon $[1, K]$, we maximize the objective function given by (2) at $k = 0$. The optimization algorithm follows a regenerative optimization scheme, i.e., it belongs to the class of open-loop feedback optimal policies [15]. However, of the decisions that are made today $(k = 0)$, only the commands $x_i(k)$ are to be implemented at $k = 1$. Thus, it is possible to use the previous optimization results as initial conditions for the next periods optimization. Note that our formulation allows a cell to be covered by multiple interdiction assets. The interdiction objective is formulated as:

$$\begin{aligned}
&\max_{\mathbf{U}} \sum_{k=1}^{K} \gamma^{(k-1)} \sum_{g \in G} PA(g,k) \\
&\qquad\qquad \bullet \left[1 - \prod_{i \in I_k} (1 - PI_i(x_i(k), g) PD(g,k))\right] \\
&\text{s.t. } \ x_i(k+1) \in R_i[x_i(k)];\ x_i(0) \text{ is given}; i \in I_k; \\
&\qquad k = 0, ..., K-1
\end{aligned} \tag{2}$$

where $\gamma (0 \leq \gamma \leq 1)$ is the discount factor. The interdiction probability $PI(xi(k), g)$ is calculated based on the centered 1-D scheme proposed in [4], which takes into consideration the vessel speed, the helicopter speed (if any), and the delay time to launch the helicopter. Following [4], the probability

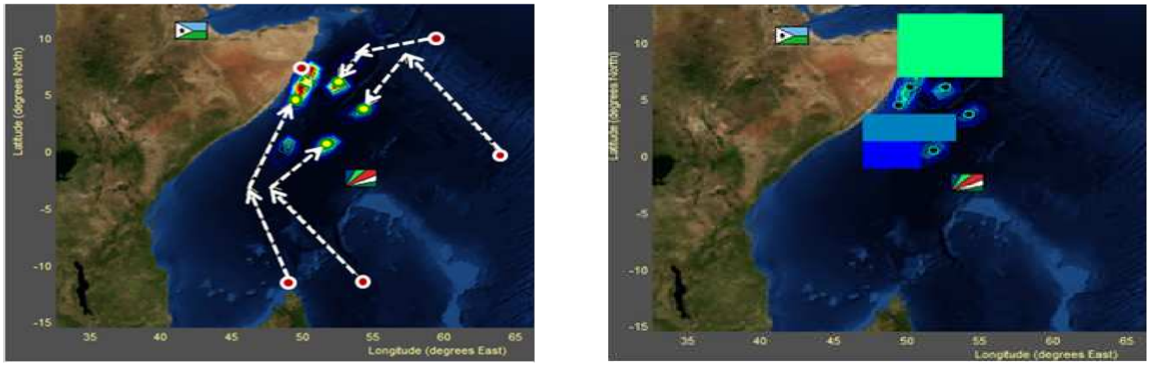

(a) (b)

Fig. 4. (a) PARS Maps and Routes of Interdiction Vessels; (b) Surveillance Search Boxes [3]

of interdicting a piracy event in cell g within a specified interdiction time τ (typically 30 minutes) is given by:

$$PI_i\left(x_i\left(k\right), g\right) = \begin{cases} \frac{2r(i,\tau)}{dist(x_i(k),g)}, r\left(i,\tau\right) < \frac{dist(x_i(k),g)}{2} \\ 1, \qquad r\left(i,\tau\right) \geq \frac{dist(x_i(k),g)}{2} \end{cases} \tag{3}$$

where $dist(x_i(k), g)$ is the Euclidean distance from cell g, where a piracy event takes place, to asset i's location $x_i(k)$, and $r(i, \tau)$ is the distance that will be covered in time τ by asset i. It is defined as:

$$r\left(i,\tau\right) = \begin{cases} v_i\tau, & \tau \leq t_i^h \\ v_i t_i^h + v_i^h\left(\tau - t_i^h\right), & \tau > t_i^h \end{cases} \tag{4}$$

where v_i is the speed of asset i, v_i^h is the speed of the helicopter operated by asset i, and t_i^h is the time to launch the helicopter (typically 10 minutes). The detection probability PD of asset i (with surface radar range $\rho(i)$), at time k is:

$$PD\left(g,k\right) = \begin{cases} 1, & dist\left(x_i\left(k\right), g\right) \leq \rho(i) \\ 0, & \text{otherwise} \end{cases} \tag{5}$$

Realistic constraints such as the weather, asset availability, asset capabilities (e.g., range, speed), and asset assignment (e.g., many sensors may need to be coordinated to obtain a better picture of the situation) are considered [3].

B. Surveillance Operations

In general, assets used for large ocean surveillance generally perform one surface search mission per day. The assets are assigned to predefined search regions in a "box," where the actual search time allowed for the asset is determined by excluding the transit time from the time interval, Δ. In this section, we formulate the surveillance problem, where the area not covered by interdiction assets is partitioned into search regions having rectangular shapes. A set S_k of available surveillance assets at time k is assigned to the partitioned regions to maximize the discounted cumulative sum of detection probability over the planning horizon, $k = 1, 2, .., K$. A search region assigned to a surveillance asset s at time k is given by the set of cells $A_s(k)$, which is a rectangular subset of cells in G. It is defined by two coordinates comprising a longitude and latitude. Each class of surveillance assets, $s \in S_k$, has different capabilities, measured in terms of the sweep width $w_s(k)$ and the speed $v_s(k)$. Note that the effective sweep width $w_s(k)$ of asset s is a function of METOC conditions in the region at a particular time k. Let the probability map of pirate presence be such that $PP(g, k)$ denotes the probability that at least one pirate is in cell g at time k. Let the amount of time spent by asset s, $s \in S_k$, in the assigned search region A_s during time step k be given by $\tau_s(k)$. Following [5][6], the probability of detection of asset s assigned to a set of cells $A_s(k)$ is given by:

$$\begin{aligned} &PD_s\left(A_s\left(k\right), k\right) \\ &= \sum_{g \in A_s(k)} PP\left(g,k\right) \times \left(1 - \exp\left(\frac{-w_s(k) * v_s(k) * \tau_s(k)}{a_c |A_s(k)|}\right)\right) \end{aligned} \tag{6}$$

where a_c is the area of a cell and $|As(k)|$ is the number of cells in the search region, $A_s(k) \subset G$. The surveillance problem can be succinctly written as follows:

$$\begin{aligned} &\max_{\{A_s(k), s \in S_k\}} \sum_{k=1}^{K} \gamma^{(k-1)} \sum_{s \in S_k} PD_s(A_s(k), k) \\ &\text{s.t.}\, A_i\left(k\right) \cap A_j\left(k\right) = \emptyset,\, (i \neq j) \in S_k \\ &\quad A_i\left(k\right) \text{has rectangular shape,} \forall i \in S_k \end{aligned} \tag{7}$$

Our technical approach to solving this NP-hard optimization problem involves decomposing it into two sequential phases by exploiting the fact that interdiction assets are substantially slower than surveillance assets. In Phase I, we solve the problem of allocating only the interdiction assets, such that regions with high cumulative probability of attack over the planning horizon are maximally covered as shown in Figure 4a. Subsequently, in Phase II, we solve the surveillance problem, where the area not covered by interdiction assets is partitioned into non-overlapping search regions (e.g., rectangular boxes) and assigned to a set of surveillance assets to maximize the cumulative detection probability over the planning horizon. In order to overcome the curse of dimensionality associated with Dynamic Programming, we developed a Gauss-Seidel

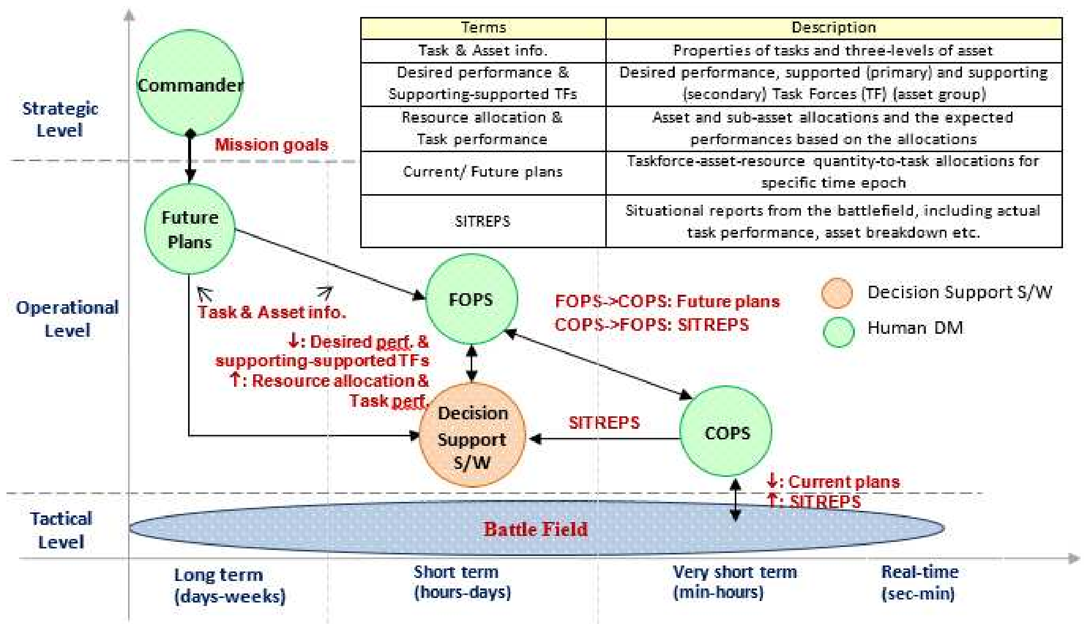

Terms	Description
Task & Asset info.	Properties of tasks and three-levels of asset
Desired performance & Supporting-supported TFs	Desired performance, supported (primary) and supporting (secondary) Task Forces (TF) (asset group)
Resource allocation & Task performance	Asset and sub-asset allocations and the expected performances based on the allocations
Current/ Future plans	Taskforce-asset-resource quantity-to-task allocations for specific time epoch
SITREPS	Situational reports from the battlefield, including actual task performance, asset breakdown etc.

Fig. 5. Maritime Operations Center Experiments [9]

algorithm, coupled with a rollout strategy, for the interdiction problem [15][16]. For the surveillance problem, we developed a partitioning algorithm followed by an asymmetric assignment algorithm for allocating assets to the partitioned regions. Once the surveillance assets are assigned to the search regions, the search path for each asset is determined based on a specific search strategy within the search region, as shown in Figure 4b [3]. In the case of counter-smuggling missions, the surveillance asset allocation is made *before* interdiction, because the surveillance assets have the flexibility to be allocated/re-allocated more frequently as compared to interdiction assets. However, exhaustive search at a particular search box may sacrifice the opportunity to detect other critical smugglers in the vicinity of the search area. Consequently, interdiction of time-critical targets requires frequent updates to case information so that appropriate decisions can be made to (re)allocate the assets. The above interdiction algorithm was recently enhanced to dynamically place the interdiction assets in the areas that are being patrolled by surveillance assets and to prioritize the interdiction areas using the projected probability of detection of the surveillance assets.

III. Distributed and Collaborative Dynamic Resource Allocation

Dynamic resource management process becomes more complex especially for Maritime Operations Centers (MOCs) where emphasis is on standardized processes and methods, centralized assessment and guidance, networked distributed planning, and decentralized execution for assessing, planning and executing missions across a range of operations [8]-[10]. The planning process, which is informed by guidance from higher headquarters along with a concomitant assessment process, should be collaborative both vertically (with higher headquarters or with lower-level subordinate task forces that "own" the individual assets) and horizontally (with other MOCs or joint components). It is essential that the planning process should be aligned with a specific organizational structure (an asset, e.g., a UAV, reports to a specific Task Force) and should consider the fact that assets are organized hierarchically (e.g., a platform or asset contains sub-platforms/sub-assets, which have *capabilities*). As MOC planning is in a distributed and collaborative setup, it is crucial to determine the key planning elements such as *who* has the expertise to make the plan (DMs who may be humans or autonomous agents), *what* needs to be planned (tasks, jobs, actions to be executed using assets or resources), *why* make the plan (desired goal or objective function), *how* to achieve the expected outcome (the assignment of assets to tasks, sequencing of activities arranged as a task graph), *where* the plan is executed (task location or mission area), *when* the plan is executed (start time and duration for each task) and *with what* facilities to make the plan (information about tasks, assets, desired objectives, decision support systems, etc.). The interactions of the above parameters include where DMs make decisions regarding the allocation (*how, when and where*) based on the information available on the assets they own and the tasks they are responsible for.

A typical operational level planning process in an abstracted MOC is shown in Figure 5. It includes the following intelligent entities: *Future Planners* (*FP*) collaboratively convert the higher-level mission goals or commander's intent into a Course of Action (COA) for the mission. This COA decomposes the

Fig. 6. VOI Analysis: Information Surfaces (Uniform, Flow and BonD)

mission goals into a graph of sub-goals or specific tasks to achieve the goals, and also includes estimated requirements and available resources to accomplish every task, and, ultimately, the commander's intent. Each sub-goal in the COA is represented as a task graph with branches and sequel options. Prior intelligence, historical and forecasted weather patterns and logistics play key roles in the development of future plans. *Future Operations* (*FOPS*) allocate assets to tasks based on the *FP*-specified COA. This allocation is made over a moving time horizon (typically a 3-day horizon, day *T*, *T+1*, *T+2*, where *T* is the current day), taking into account dynamically evolving intelligence, logistics and weather information provided by weather reach back cells. *Current Operations* (*COPS*) monitor the ongoing activities on day *T* and provide feedback to *FOPS* and *FP* in the form of situation reports (SITREPS) on emerging tasks and their requirements, task outcomes, changes in task requirements, asset (and network) status, and evolving intelligence. *Multilevel planning agents* provide information and decision support to assist *FOPS* planners in evaluating and optimizing the asset-to-task allocation at several levels. At the Task Force (*TF*) level, the agents suggest different supporting-supported options across a number of interdependent tasks in competing task graphs (representing different missions), taking into account uncertainty in weather forecasts, intelligence, asset status and network status. At the platform and warfare area levels, the agents optimize the sub-platform-to-task allocation. Other agents compute mission context-dependent value of information and decisions and manage the flow of information among decision makers (DMs) [7]-[9].

The MOC research problems can be broadly categorized as: 1) Modeling, formulating, and solving a domain-independent planning problem of allocating asset packages to a set of interdependent tasks, where each task is specified by a vector of resource requirements, and each asset is specified by a vector of resource capabilities; 2) Adapting the planning process to a distributed and collaborative environment, where multiple DMs having ownership of disjoint assets, have different information and expertise collaborate to achieve a mission objective; 3) Integrating realistic multi-level organized assets into the planning problem formulation.

In order to address the above issues, we designed, developed, and deployed a library of dynamic resource allocation algorithms including: *Asset package selection-based scheduling* algorithm solves the planning problem by specifying asset packages on a task-by-task basis. The weighted length algorithm (WL) is applied to prioritize the tasks in the task graph based on their processing time and its position in the task graph. For each task, we obtain *m*-best asset packages (*m* is a user-specified parameter). A rollout strategy is employed to evaluate the impact of allocating an asset package to a particular task on the execution accuracies of unassigned tasks in the task list, thereby improving the WL solution. The solution is further improved by using a pair-wise exchange (PWE) heuristic that considers all possible task assignment sequences obtained by exchanging the task at the current place in the assignment sequence with some other task, while not violating the precedence constraints [7].

Blackboard-based collaborative planning framework enables information sharing and resource transfer among DMs to achieve a certain mission objective. The DMs use the Asset Package Selection module to allocate assets to their own tasks; we call this the *intra-agent* module. The *inter-agent* module then interacts with the blackboard by sharing asset and task information as well as asset prices that quantify the importance of an asset to each individual DM. Another important aspect of our architecture is a model of the DMs cooperative behavior in terms of the relative priority they give to the other DMs tasks. We conducted computational experiments to investigate how various cooperative behaviors affect mission performance (measured in terms of average task accuracy), communication cost (in terms of shared assets and tasks) and workload (measured in terms of the number of tasks they are responsible for and the number of tasks they are assigned in the final plan). Mission performance was compared for five cases: "no collaboration" among DMs, "self-interested" DMs who are only concerned with their own tasks, "teamwork" where each DM treats every task as being equally important, "benevolent" where each DM gives higher priority to other DMs' tasks than his own, and the "centralized" case. The results indicate that as the cooperative

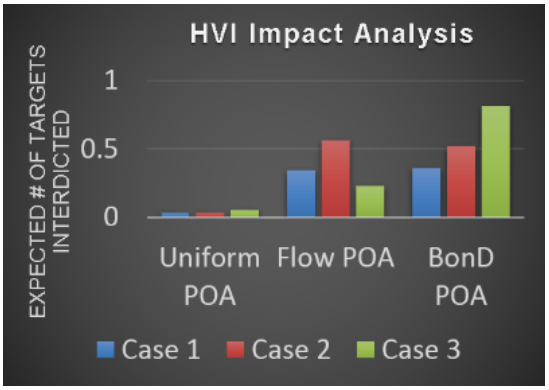

Fig. 7. Value of Information: Performance Analysis

behavior increases, DMs send a greater share of their tasks and assets to be allocated by the Coordinator, i.e., it becomes a centralized planning problem requiring greater communication and coordination. A good trade-off can be found between teamwork and benevolent cooperative behaviors in order to maintain a balanced workload, reasonable communication cost and near-optimal planning performance [8].

Further, we extended the asset package selection-based planning algorithm to a *multi-level asset allocation problem* to minimize the overall difference between a human-specified task accuracy performance criteria and the expected performance based on how well the assigned resources match the required resources, subject to a number of real-world planning constraints. Due to hierarchical decomposition of assets and complex constraints, the problem size grows dramatically as more tasks are to be planned for. We developed two methods: (1) a Lagrangian Relaxation-based planning algorithm that decomposes the problem into two solvable sub-problems, viz., a linear programming sub-problem with discrete variables and a nonlinear programming problem with continuous variables; (2) a Dynamic List planning algorithm that iteratively assigns the most preferred asset to the most demanding task until the tasks' accuracies are as close to the desired accuracies as possible. The experimental results in [9] demonstrate that the Dynamic List Planning method is a highly-efficient near-optimal solution, while the Lagrangian relaxation method is a highly accurate, but somewhat slower algorithm. The proposed algorithms are verified using realistic MOC planning scenarios, by providing a comparative evaluation of the performance measures of the two proposed methods, and investigating the value of information via human-in-the-loop experiments [9].

IV. Value of Information

Examining decision quality from different perspectives (i.e., mission goals, impact of the environment, asset status and task status) will enable mission relevant high value information (HVI) to be identified, acquired, and delivered to the right DM at the right time to achieve high quality decisions. Paucity of information results in poor decisions due to not having enough situation relevant data; on the other hand, having too much information will distract and overburden the DM resulting in poor decision quality. Finding HVI that maximizes decision quality will enable decision support systems to recommend effective COAs for mission success. When information is accurately valued, extracted, and prioritized, it pre-stages decision relevant information, alleviates the bandwidth limitations in distributed, intermittent and low bandwidth (DIL) environments, thereby promoting mission success.

Quantification of the value of information (VOI) allows acquisition and integration of the right data from the right sources in the right context to the right DM (human) at the right time for the right purpose (known as the 6R concept [1]), providing a means to define and value unique information related to context and decision models. In this section, we discuss a variety of methods to quantify the value of information in a DM-understandable way. The methods include both tangible and statistical variations that are based upon situation-driven mission objectives and different types of information gains or distances, respectively. Given a mission objective and situation query, a VOI sensitivity analysis can be conducted to offer insights as to how degradation and quality of information may impact mission outcomes and situation awareness. Different types of information pertaining to mission, environment, assets, and pending tasks can be analyzed and in doing so, their value can be quantified based on tangible mission-relevant metrics. Measures of information value based on Bayesian diagnosticity, impact, information gain, and other Bayesian Optimal Experimental Design framework theories [11] are used widely outside of operations research and elaborated on in [13]. Statistical entropy-based computations and distances such as the aforementioned and including pre-posterior analysis, utility, and Kullback-Leibler Divergence serve as metrics to prioritize information, connecting the calculated number with a piece of information deemed important or irrelevant to the current context. While these statistics-based methods do not have a tangible unit, when computed for multiple pieces of information, based on comparison, prioritization is feasible [11].

In addition to the tangible and entropy-based metrics, we can also evaluate the value of information by conducting sensitivity analysis of mission metrics with respect to availability and non-availability of information as shown in Figure 6. Finding the peak of the decision quality curve is integral to exploring how high quality, effective COAs can be suggested via proactive decision support tools. A prototype version of a "wrapper-based" approach to information valuation has been explored in the context of counter-smuggling, where we examined the impact of three different degrees of uncertainty propagated onto PoA surfaces and the concomitant expected decision quality [11]. As shown in Figure 7, providing PoA surfaces in the form of historical routes of smugglers provided 884% improvement in decision quality (as measured by expected number of targets interdicted) in comparison to a uniform PoA surface, while the PoA specific to the cases (termed Battlespace on Demand (BonD) PoA) provided 50% further improvement in relation to the historical routes and 1445%

improvement overall. Thus, case-based high-valued contextual information in the form of PoAs can substantially improve mission success. We plan to implement similar context-driven analyses in other mission contexts (e.g., ASW, ISR, UAV) by considering the uncertainty in PoAs themselves ("second order uncertainty or ambiguity" [14]) and evaluating the sensitivity of HVI with respect to this uncertainty [11].

V. Conclusion

In this paper, we briefly discussed dynamic asset allocation algorithms for surveillance and interdiction operations in the context of counter-smuggling and counter-piracy missions. The dynamic surveillance and interdiction asset allocation problems are NP-hard. In order to overcome the curse of dimensionality, we proposed the method of successive displacements and rollout concepts for solving the interdiction problem. For the surveillance problem, we proposed a partitioning algorithm, where each region is grown independently subject to the regions shape constraints. We also presented distributed and collaborative resource management algorithms, which are capable of interacting with human decision makers in providing asset package selection, blackboard-based collaborative planning and multi-level dynamic asset allocation in dynamic and uncertain environments. Additionally, techniques and methods to quantify the value of information were briefly discussed.

The proposed algorithms facilitate efficient utilization of assets at the operational level by providing intelligent courses of action to the appropriate commanders in a timely manner. Additionally, it facilitates the DMs in promptly understanding and envisioning the current and projected mission context, while allowing them adequate time to make appropriate decisions by taking into account the concomitant uncertainties, and unknown risks stemming from the specific context via networking, collaboration, distributed execution, and resource sharing within the mission environment. Our future work will focus on exploring approximate dynamic programming (ADP) [17] techniques for dynamically coordinating surveillance and interdiction assets in a dynamic and uncertain environment and planning for unexpected scenarios, e.g., asset breakdown, pop-up threats, etc.

Acknowledgment

The authors would like to thank NRL-MRY, SPAWAR, and NPS for their valuable contributions in designing and detailing the piracy/drug trafficking scenarios and the concept of operations embedded within our decision support software tool. This work was supported by the U.S. Office of Naval Research under contract # N00014-12-1-0238 and # N00014-09-1-0062.

References

[1] A. Smirnov, "Context-Driven Decision Making in Network-Centric Operations: Agent-Based Intelligent Support," in CKM Workshop, 2006, vol. 7, no. 812.

[2] D. McCarren, "Navy NUOPC Requirements Battlespace on Demand", Earth System Prediction Capability Workshop, Colorado, 2010.

[3] W. An, D. F.M. Ayala, D. Sidoti, M. Mishra, X. Han, K. R. Pattipati, E. Regnier, D. L. Kleinman and J. Hansen, "Dynamic Asset Allocation Approaches for Counter-piracy Operations," *15th International Conference on Information Fusion*, Singapore, July 9-12, 2012.

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

[5] B.O. Koopman, "Search and Screening: General Principles with Historical Applications," *Pergamon Press*, New York, 1980.

[6] U.S. Coast Guard Research and Development, "Review of search theory: advances and applications to search and rescue decision support," *Report*, 2001.

[7] X. Han, H. Bui, S. Mandal, K.R. Pattipati and D.L. Kleinman, "Optimization-based Decision Support Software for Team-in-the-loop Experiment MOC-1: Asset Package Selection and Planning," *IEEE Trans. on SMC: Part A Systems*, Vol. 43, No. 2, pp. 237-251, March 2013.

[8] X. Han, S. Mandal, K. R. Pattipati, D. L. Kleinman, and M. Mishra, "An Optimization-based Multi-Agent Distributed Planning Algorithm: A Blackboard-Based Collaborative Framework," *IEEE Trans. On SMC: Part A Systems*, Vol. 44, No. 6, pp. 673-686, June 2014.

[9] X. Han, M. Mishra, S. Mandal, H. Bui, D.F. Martinez, D. Sidoti, K.R. Pattipati and D.L. Kleinman, "Optimization-based Decision Support Software for Team-in-the-Loop Experiment: Multi-level Asset Allocation," *IEEE Transactions on SMC Systems*, Vol. 44, No. 8, pp. 1098-1112, August 2014.

[10] S. G. Hutchins, W. G. Kemple, D. L. Kleinman, S. Miller, K. Pfeiffer, S. Weil, Z. Horn, M. Puglisi, and E. Entin, "Maritime headquarters with maritime operations center: A research agenda for experimentation," In *Proceedings of the 14th International Command and Control Research and Technology Symposium (ICCRTS)*, June 2009.

[11] D. Sidoti, D. Ayala, X. Han, M. Mishra, S. Sankavaram, W. An, K. R. Pattipati, and D. Kleinman, "Evaluating the Value of Information in the Presence of High Uncertainty," in *18th ICCRTS: C2 in Underdeveloped, Degraded, and Denied Operational Environments*, 2013.

[12] D. Sidoti, X. Han, D. F.M. Ayala, M. Mishra, S. Sankavaram, W. An, K.R. Pattipati, "An Approximate Dynamic Programming Approach for Context-Aware Dynamic Asset Allocation for Multiple Targets" to be submitted to *IEEE Transactions on Systems, Man and Cybernetics*.

[13] J. Nelson, "Towards a rational theory of human information acquisition," *The probabilistic mind: Prospects for rational models of cognition*, pp. 143163, 2008.

[14] H. J. Einhorn, and R. M. Hogarth, "Ambiguity and uncertainty in probabilistic inference," *Psychological Review*, Vol. 92, 4, 1985.

[15] D. P. Bertsekas, *Dynamic Programming and Optimal Control/Approximate Dynamic Programming*, Athena Scientific, 2007.

[16] D. P. Bertsekas, *Nonlinear Programming*, Athena Scientific, Belmont, MA, 2003.

[17] W. B. Powell, *Approximate Dynamic Programming: Solving the Curses of Dimensionality*, Wiley-InterScience, Hoboken, NJ, 2007.

[18] D.M.F. Ayala, D. Sidoti, M. Mishra, X. Han, and K. R. Pattipati, "Context-based Models to Overcome Operational Challenges in Maritime Security", accepted to *IEEE Symposium on Technologies for Homeland Security*, April, 2015.

[19] http://www.airforcetimes.com/article/20100402/NEWS/4020326/UAV-crash-blamed-distraction-battle