Decentralized Task Allocation with Coupled Constraints in Complex Missions

Andrew K. Whitten, Han-Lim Choi, Luke B. Johnson, and Jonathan P. How

Abstract— This paper presents a decentralized algorithm that can create feasible assignments for a network of autonomous agents in the presence of coupled constraints. The coupled constraints are introduced to address complex mission characteristics that include assignment relationships, where the value of a task is conditioned on whether or not another task has been assigned, and temporal relationships, where the value of a task is conditioned on when it is performed relative to other tasks. The new algorithm is developed as an extension to the consensus-based bundle algorithm (CBBA), introducing the notion of pessimistic or optimistic bidding strategies and the relative timing constraints between tasks. This extension, called Coupled-constraint CBBA (CCBBA), is compared to the baseline in a complex mission simulation and is found to outperform the baseline, particularly for task-rich scenarios.

I. Introduction

Decentralized task planning for a team of networked agents (e.g., unmanned aerial vehicles (UAVs)) has been investigated by many researchers [1]–[13]. Ideally, the communication link between all elements of the autonomous decision system (command station, autonomous vehicles, manned vehicles, etc.) is high bandwidth, low latency, low cost, and highly reliable. However, even the most modern communication infrastructures do not posses all of these characteristics. If the inter-agent communication mechanism has a more favorable combination of these characteristics compared to agent-to-base communication, then a decentralized planning architecture offers performance and robustness advantages. In particular, response times to changes in situational awareness can be significantly faster via decentralized control than those achieved under a purely centralized planner. In addition, decentralized planning schemes are well-suited for situations where the information needed for decision making is local with respect to the network diameter. This is particularly noticeable in the task assignment consensus problem where agents near each other will require the most communication to resolve task assignment conflicts, whereas agents that are spatially separated are less likely to choose the same tasks. Decentralized algorithms with strong inter-agent communication should support this scenario more efficiently than centralized approaches.

One decentralized tasking approach the authors have recently developed is the consensus-based bundle algorithm (CBBA) [12], which is a market-based distributed agreement protocol upon the winning agents and associated winning

A. K. Whitten, L.B. Johnson, and J. P. How was/are with the Dept. of Aeronautics and Astronautics, MIT, Cambridge, MA. {awhitten, lbj, jhow}@mit.edu.

H.-L. Choi is with the Div. of Aerospace Engineering, KAIST, Korea. hanlimc@kaist.ac.kr

scores. By enforcing agreement upon the solution rather than the information set, CBBA was shown to account for inconsistent information among the networked agents, guaranteeing a conflict-free assignment for all the agents in the network. CBBA is a polynomial-time algorithm that scales well with the size of the network and/or the number of tasks (or equivalently, the length of the planning horizon). From the design perspective, various design objectives, agent models, and constraints can be incorporated by defining appropriate scoring functions. If the resulting scoring scheme satisfies a property called *diminishing marginal gain* (DMG), a provably good feasible solution is guaranteed (50% minimum performance in theory; over 90% in simulations). Furthermore, a recent extension to CBBA [13] has enabled incorporation of heterogeneity in the agent capabilities and time windows of validity of tasks, which significantly enriches the mission characteristics that can be handled.

However, this recent progress is still not enough to address task planning for complex missions wherein tasks are coupled both spatially and temporally. The main source of such coupling is the set of rules of engagement, which can often be described as logical constraints and temporal precedence relationships. This paper particularly focuses on embedding these coupled constraints in the CBBA-based decentralized tasking framework. While a preliminary approach to achieve this functionality was presented in [14] for specific types of constraints described by cooperation requirements and preference, this paper provides a more systematic methodology that can handle a much broader class of coupled constraints. Numerical experiments on cooperative sensing & neutralization mission of uninhabited aerial vehicles (UAVs) verify the performance advantage of the proposed extension compared to the baseline CBBA approach.

II. Task Allocation with Constraints

Consider a bounded environment containing a network of N_u autonomous agents, and a set of N_t tasks. For notational convenience, both agents and tasks are given integer indices. The set of all agent indices is denoted $\mathcal{I} \triangleq \{1, \ldots, N_u\}$, while the set of all task indices is denoted $\mathcal{J} \triangleq \{1, \ldots, N_t\}$.

Each task, $j \in \mathcal{J}$, is given a score function, $S_j(a_j, t_j) \in \mathbb{R}_+$ which represents the value that task adds to the mission as a function of $a_j \in \mathcal{I} \cup \{\emptyset\}$, the index of the agent assigned to task j, and $t_j \in \mathbb{R}_{\geq 0} \cup \{\emptyset\}$ the time the agent plans to arrive at the j-th task location. The symbol $\emptyset$ denotes a null entry in the corresponding set. The goal of the task planner is to assign agents to tasks in such a way that the cumulative score over all tasks in the task set is maximized. The task

assignment problem at a given mission time has an objective function expressed as

$$\max \sum_{j=1}^{N_t} S_j(a_j, t_j) \tag{1}$$

Note that in this formulation, there exists no requirement to assign an agent to every task, and in fact doing so may be infeasible. However, a task only earns a positive score if an agent is assigned to it, $S_j(a_j, t_j) > 0$ only if $a_j \neq \emptyset$. Therefore, the task planner must determine an appropriate subset of tasks to assign, and select the agent, and visit time for each assigned task. The assignment is thus completely determined by selecting a_j and t_j for each task $j \in \mathcal{J}$.

The problem is subject to two types of constraints: 1) agent-specific constraints, and 2) coupled constraints. Agent-specific constraints affect the options available to a given agent i independently of decisions made with respect to all other agents, and include:

1) *Capability constraints*: Each task has a requirement, and each agent has a set of capabilities. Compatibility matrix, $\mathcal{M}$ captures each agent's fitness for performing the task. $\mathcal{M}_{ij} > 0$ if agent i is allowed to perform task j, and zero otherwise.
2) *Time window of validity*: Each task j, has a time window of validity given by $[\tau_j^{\text{start}}, \tau_j^{\text{end}}]$. The task must be started on the interval of time specified by the time window. If a task is not time sensitive, it is given a time window of $[0, \infty)$.
3) *Vehicle dynamics*: Each agent has a set of dynamics that impose a set of constraints on that agent. Agent i has a maximum speed, $v_i^{\max}$, a maximum acceleration, $\dot{v}_i^{\max}$, and a minimum turn radius $r_i^{\min}$. These parameters determine the minimum amount of time required between each task the agent performs.
4) *Fuel constraints*: At a given time t, agent i has remaining fuel mass, $m_{\text{fuel}}^i(t)$, and has nominal fuel consumption rate, $\dot{m}_{\text{fuel}}^i$. These parameters determine the maximum remaining time aloft for agent i.

Coupled constraints include any situation where the decisions regarding one agent alter the options available to another agent. Often, in a complex mission, there are a variety of such constraints due to rules of engagement or simply the nature of the tasks involved. Coupled constraints include:

• *Assignment constraints* deal with the relationship between tasks with respect to which combination of assignments are permitted. Examples:

1) *Conflict-freeness:* Each task may have at most one agent assigned to it. (Note that this is a simplifying assumption: mission activities requiring or benefiting from multiple agents may be represented by multiple instances of the same task.)
2) *Unilateral Dependency:* A task, A is dependent on another task B, but task B is not dependent on A. In other words, task A cannot be assigned unless B is also assigned, but B may be assigned independently of task A.
3) *Mutual Dependency:* A task A is dependent on another task B, and task B is dependent on task A. Task A and task B must be assigned together or not at all.
4) *Mutually Exclusive:* A task A cannot be assigned if another task B is assigned, and task B cannot be assigned if task A is assigned.

• *Temporal constraints* include any specified relationship between the chosen visit times for a subset of tasks. Common temporal constraints include, but are not limited to:

1) *Simultaneous:* task A and B must begin at the same time.
2) *Before:* task A must end before task B begins.
3) *After:* task A must begin after task B ends.
4) *During:* task A must begin while task B is in progress.
5) *Not during:* task A must either end before task B begins, or begin after task B ends.
6) *Between:* task A must begin after task B ends and end before task C begins.

III. CBBA: THE BASELINE TASKING ALGORITHM

The following summarizes CBBA first developed in [12] and recently extended in [13].

A. Phase I: Task Selection

In the first phase of CBBA, *task selection*, each agent iteratively adds tasks to its bundle. The selection process is *sequentially greedy*; given an initial bundle, $\mathbf{b}_i$, agent i adds the task that will result in the largest marginal increase in score, and repeats the process until the bundle is full, or no more tasks can be added. Each time agent i adds a task j to its bundle, it enters its own index into the winning agent list, $z_{ij} \leftarrow i$, and its corresponding bid into the winning bids list, $y_{ij} \leftarrow S_j(i, t_{ij})$. Agent i is only allowed to add task j to its bundle if agent i can place a higher bid than the current highest bidder, y_{ij}.

B. Phase II: Consensus

After all agents complete a round of the task selection phase, agents communicate with each other to resolve conflicting assignments within the team. After receiving information from neighboring agents about the winning agents and corresponding winning bids, each agent can determine if it has been outbid for any task in its bundle. Since the bundle building recursion (Section III-A) depends at each iteration upon the tasks in the bundle up to that point, if an agent is outbid for a task, it must release it and all subsequent tasks from its bundle. For each message (consisting of bidding information on a task) that is passed between a sender k and a receiver i, a set of actions is executed by agent i to update its information vectors using the received information. These actions involve comparing its winning agent list $\mathbf{z}_i$, the winning bids list $\mathbf{y}_i$, and the time stamp vector $\mathbf{s}_i$ to those of agent k to determine which agent's information is the most up-to-date for each task.

Fig. 1. Example of task set partitioning

C. Properties

The baseline CBBA is capable of handling non-coupled constraints including capability constraints, time windows, and basic vehicle dynamics such as top speed. It is also proven that the baseline algorithm is also capable of guaranteeing conflict-free assignments, assuming a strongly connected network, which is one type of coupled constraint.

However, it is not able to account for the other types of coupling which are common in complex missions. This paper describes the necessary machinery for enforcing the following types of coupled constraints within the CBBA framework: 1) unilateral dependency constraints, 2) mutual dependency constraints, 3) mutual exclusions, and 4) temporal constraints.

IV. COUPLED-CONSTRAINT CBBA (CCBBA)

A. Task Set Partitioning

For book-keeping purposes, the task set is partitioned into sub-groups of tasks that share coupled constraints. Each of these sub-groups is called an *activity*. For notational convenience, each task in an activity is referred to as an *element* of that activity, and is given an index, $q \in \mathbb{Z}_+$. Each task in the task set can be uniquely identified by its activity number and element index, and the notation j_q is used to indicate the task associated with the q^{th} element of activity j. The set of all tasks in the task set is still denoted $\mathcal{J}$. The set of all activities in the set is denoted $\mathcal{A}$, while the set of all elements in an activity j is denoted $\mathcal{A}_j$. Each pair of elements in an activity may have a coupled constraint between them but does not have to. However, by construction, coupled constraints do not exist between tasks belonging to different activities with the exception of the conflict-freeness constraint.

An example is shown in Fig. 1: task B may not be assigned unless task A is assigned, but task A may be assigned independently of task B. Task C may not be assigned unless task D is assigned, and task D may not be assigned unless task C is assigned. The task pair (C, E) may not both be assigned, and the task pair (D, E) may not both be assigned. This particular task set can be grouped into two activities, with two and three elements respectively. The task set $\mathcal{J} = \{A, B, C, D, E\}$ is now written $\mathcal{J} = \{1_1, 1_2, 2_1, 2_2, 2_3\}$.

TABLE I
CODE FOR DEPENDENCY MATRIX ENTRY $\mathcal{D}_{qu}^j$

1	u depends on q
0	u may be assigned independently of q
-1	q and u are mutually exclusive
$a > 1$	u requires either q or another element with the same code, a. Entries are used sequentially: 3 is not used unless 2 is used, etc.

For each activity $j \in \mathcal{A}$, the constraint structure can be compactly written in the form of a *Dependency Matrix*, $\mathcal{D}^j$ where the row column entry, (q, u), denoted as $\mathcal{D}_{qu}^j$, describes the relationship between the q^{th} element of activity j and the u^{th} element of activity j. The notation for encoding the constraints is provided in Table I. The diagonal entries of $\mathcal{D}^j$ are defined to be 0 for all activities.

Each task in the task set is given a *bidding strategy* which represents how the bidding protocol is handled; the bidding strategy for each task $j_q \in \mathcal{J}$ is determined from the dependency matrix as

$$\text{bidStrat}_{j_q} = \begin{cases} \text{optimistic} & \text{if } O_{j_q} \neq \emptyset \\ \text{pessimistic} & \text{otherwise} \end{cases}$$

where $O_{j_q} \triangleq \{u : \mathcal{D}_{uq}^j \geq 1 \wedge \mathcal{D}_{qu}^j = 1\}$.

B. Pessimistic Bidding Strategy in Phase I

In the baseline CBBA, at the beginning of the task selection phase, an agent i calculates the marginal score for each of the tasks in the task set that are not in their bundle. In CCBBA, a step is added before this where an agent determines which tasks it is allowed to bid on.

Task j_q is given a pessimistic bidding strategy (PBS) if it is only dependent upon tasks which do not depend on j_q. The PBS specifies that agent i is allowed to bid on task j_q only if all of the dependency constraints for task j_q are satisfied. An indicator $\text{canBid}_i(j_q)$ is used to keep track of this permission:

$$\text{canBid}_i(j_q) = \begin{cases} \text{true} & \text{if } n^{\text{sat}}(j_q|\mathbf{z}_i) = N_{\text{req}}(j_q) \\ \text{false} & \text{otherwise} \end{cases}$$

where $N_{\text{req}}(j_q)$ and $n^{\text{sat}}(j_q|\mathbf{z}_i)$ denote the total number of constraints required to be satisfied, and the number of satisfied constraints according to the current knowledge, respectively. For the case where $\mathcal{D}_{uq}^j \in \{-1, 0, 1\}, \forall u, q$, these quantities are given by

$$N_{\text{req}}(j_q) = \sum_{u=1}^{|\mathcal{A}_j|} \mathbb{I}\left(\mathcal{D}_{uq}^j = 1\right)$$
$$n^{\text{sat}}(j_q|\mathbf{z}_i) = \sum_{u=1}^{|\mathcal{A}_j|} \mathbb{I}\left((z_{i(j_u)} \neq \emptyset) \wedge (\mathcal{D}_{uq}^j = 1)\right),$$

while their expressions for more sophisticated cases can be found in [**?**, Section 3.3]. If $\text{canBid}_i(j_q)$ is true, then the marginal score for j_q is calculated to determine if this task should be added to agent i's bundle. If $\text{canBid}_i(j_q)$ is false, then the marginal score for j_q is set to 0.

C. Optimistic Bidding Strategy in Phase I

Tasks which are mutually dependent on at least one other task are given an optimistic bidding strategy. Consider a task j_q with an optimistic bidding strategy. An agent who can benefit from adding task j_q to its bundle may do so,

even if the number of satisfied constraints is smaller than the number required. The agent then keeps track of the number of iterations that pass where at least one of the tasks which j_q depends on remains unassigned. If too many iterations pass, the agent releases task j_q from its bundle. The task is released with the assumption that the other agents in the network are not interested in selecting the tasks j_q depends on, because other tasks are more valuable to these agents. To prevent an agent from repeatedly bidding on a task only to release it once it discovers no other agents are interested in performing the tasks which it depends on, the number of attempts is limited. Once an agent has run out of attempts on a particular task, they are no longer permitted to bid on that task unless all of the required constraints are satisfied. Thus, the agent has transitioned to a pessimistic bidding strategy for this task. To formalize the technique, the following definitions are introduced:

a) Number of Iterations in Constraint Violation: $\boldsymbol{\nu}_i \triangleq \{\nu_{i1}, \ldots, \nu_{iN_t}\}$ is the list which keeps track of the number of CBBA iterations which have passed with agent i violating a constraint. Element notation $\nu_{i(j_q)}$ is used to indicate the number of iterations agent i has been winning element q of activity j while at least one other element in j is unassigned.

b) Permission to Bid Solo: $\mathbf{w}_i^{\text{solo}} \triangleq \{w_{i1}^{\text{solo}}, \ldots, w_{iN_t}^{\text{solo}}\}$ indicates which elements agent i is allowed to bid on as the first agent. The list is initialized to contain positive integers representing the number of attempts an agent is given to win the particular element. If $w_{i(j_q)}^{\text{solo}} > 0$, agent i is permitted to bid on task j_q even if no other elements of j have a winning agent. If a task j_q is released due to a timeout, or a timing constraint violation, $w_{i(j_q)}^{\text{solo}}$ is decremented by one.

c) Permission to Bid Any: $\mathbf{w}_i^{\text{any}} \triangleq \{w_{i1}^{\text{any}}, \ldots, w_{iN_t}^{\text{any}}\}$ indicates which tasks agent i is allowed to bid on given that at least one of the dependency constraints is satisfied for that task. The array is initialized to contain positive integers in a similar manner to $\mathbf{w}_i^{\text{solo}}$, except the initial values are typically larger in magnitude. If $w_{i(j_q)}^{\text{any}} > 0$, and any other element of j has a winner, agent i is permitted to bid on task j_q. If an element j_q is released due to a timeout, or a timing constraint violation, $w_{i(j_q)}^{\text{any}}$ is decremented by one.

d) Timeout Parameter: For each element q of an activity j, a *timeout* parameter, o_{j_q} is defined. At each iteration, agent i increments $\nu_{i(j_q)}$ for each task, j_q, for which agent i is the winner, but the other elements belonging to activity j are not filled. If at any time, $\nu_{i(j_q)}$ exceeds o_{j_q}, the task j_q must be released from $\mathbf{b}_i$, and the values $w_{i(j_q)}^{\text{solo}}$ and $w_{i(j_q)}^{\text{any}}$, are each decremented by 1.

Whether or not an agent i is allowed to bid on task j_q is determined by:

$$\text{canBid}_i(j_q) = \big(w_{i(j_q)}^{\text{any}} > 0 \wedge n^{\text{sat}}(j_q | \mathbf{z}_i) > 0\big)$$
$$\vee \big(w_{i(j_q)}^{\text{solo}} > 0\big) \vee \big(n^{\text{sat}}(j_q | \mathbf{z}_i) = N_{\text{req}}(j_q)\big).$$

If $w_{i(j_q)}^{\text{solo}} > 0$, then the agent has not yet exhausted its attempts to acquire the proper number of partnering agents for task j_q, and so it is allowed to bid even if 0 dependency constraints are satisfied. If $w_{i(j_q)}^{\text{solo}} = 0$, but $w_{i(j_q)}^{\text{any}} > 0$,

then the agent has not yet exhausted its attempts to acquire the proper number of partnering agents, but at least one dependency constraint must be met in order for the agent to bid on task j_q. If both $w_{i(j_q)}^{\text{solo}} = 0$ and $w_{i(j_q)}^{\text{any}} = 0$, then agent i has exhausted its attempts to bid optimistically on task j_q, and may only bid if all of the dependency constraints for task j_q are satisfied.

Note that appropriate selection of o_{j_q} and appropriate initialization of $\mathbf{w}_i^{\text{solo}}$ and $\mathbf{w}_i^{\text{any}}$ are needed to achieve good performance and convergence properties. Selecting o_{j_q} too small may lead to performance degradation, discouraging agents to bid on coupled tasks, while choosing it too large may lead to long convergence time. Selecting a higher initial value for $w_{i(j_q)}^{\text{solo}}$ gives agent i more attempts to bid on task j_q, however may adversely affect the runtime of the algorithm. Selecting too small of a value may reduce performance due to conservatism in bidding on coupled tasks.

D. Mutual Exclusions in Phase I

In the baseline CBBA, an agent must be able to outbid the current winner of a task in order to be eligible to bid on that task. In Coupled-Constraint CBBA, to bid on task j_q, the marginal score for j_q must be greater than all of the tasks which are mutexed with task j_q. Therefore, task j_q is eligible to be bid on by agent i, if

$$\big(c_{i(j_q)} > y_{i(j_u)} \vee \mathcal{D}_{uq}^j \neq -1\big) = \texttt{true}, \ \forall u \in \mathcal{A}_j.$$

E. Temporal Constraints in Phase I

Often in complex missions, there exist timing constraints between some of the tasks in the task set. CCBBA allows a timing relationship to be specified between any pair of tasks belonging to the same activity. The pairwise timing constraints can be written compactly in the form of a *temporal constraint matrix*, $\mathcal{T}^j \in \mathbb{R}^{|\mathcal{A}_j| \times |\mathcal{A}_j|}$. The (q, u) entry of $\mathcal{T}^j$, denoted as $\mathcal{T}_{qu}^j$, specifies the maximum amount of time task j_q can begin after task j_u begins. Thus, task j_u must begin at least $\mathcal{T}_{qu}^j$ before task j_q begins and at most $\mathcal{T}_{uq}^j$ after task j_q begins, which can be interpreted as the *relative time window* imposed on task j_u by task j_q. There exists no sign restriction on the entries of the temporal constraint matrix, which means that the relative time window a task j_q imposes on another task j_u does not have to contain the start time for task j_q. This makes it possible to specify constraints like *before*, or *after* mentioned in Section II. If no timing constraint exists between task j_q and j_u, then $\mathcal{T}_{qu}^j = \mathcal{T}_{uq}^j = \infty$. The diagonal entries of matrix $\mathcal{T}^j$ are 0 by definition.

To satisfy coupled timing constraints, it is necessary that each agent be aware of the scheduled times of arrival of each of the tasks which have winners. Therefore, a new information list is introduced, $\mathbf{tz}_i \triangleq \{tz_{i1}, \ldots, tz_{i(N_t)}\}$. Entries are denoted $tz_{i(j_q)}$ for $j_q \in \mathcal{J}$, and the list keeps track of the arrival times for each task in $\mathbf{z}_i$ that has a winner. $tz_{i(j_q)} = \emptyset$ if task j_q has no winner according to agent i.

Agent i must calculate the time interval that is valid for each task under consideration to determine if that task is feasible given their current path, and to determine the

marginal score for that task. Agent i considering task j_q calculates the permissible start time by intersecting the original time window for task j_q with each of the coupled timing constraints imposed by the agents winning each of the tasks upon which task j_q depends, according to its knowledge, i.e., $\mathbf{z}_i$ and $\mathbf{tz}_i$.

Agent i needs to compute the interval $[\tau_{j_q}^{\min}, \tau_{j_q}^{\max}]$ where $\tau_{j_q}^{\min}$ is the minimum allowed start time for task j_q given the current knowledge of agent i and $\tau_{j_q}^{\max}$ is the maximum allowed start time for task j_q given the current knowledge of agent i. Recall that the original time window is given by $[\tau_{j_q}^{\mathrm{start}}, \tau_{j_q}^{\mathrm{end}}]$.

The minimum admissible arrival time, $\tau_{j_q}^{\min}$ is calculated by

$$\tau_{j_q}^{\min} = \max\left(\tau_{j_q}^{\mathrm{start}}, \max_{u \in \{1, \ldots, |\mathcal{A}_j|\} \setminus \{q\}} t_{\mathrm{const}}^{\min}(j_u, j_q | \mathbf{z}_i),\right)$$

where

$$t_{\mathrm{const}}^{\min}(j_u, j_q | \mathbf{z}_i) = \begin{cases} tz_{iu} - \mathcal{T}_{uq}^j & \text{if } z_{i(j_u)} \neq \emptyset \wedge \mathcal{D}_{uq}^j > 0 \\ -\infty & \text{otherwise.} \end{cases}$$

is the constraint imposed on the start time of task j_q by the agent winning task j_u. Notice that the constraint is only active if agent i believes there is an agent winning task j_u, and j_q is dependent on j_u. $\tau_{j_q}^{\min}$ is taken to be the tightest active constraint including the original time window. Similarly, $\tau_{j_q}^{\max}$ is given by

$$\tau_{j_q}^{\max} = \min\left(\tau_{j_q}^{\mathrm{end}}, \min_{u \in \{1, \ldots, |\mathcal{A}_j|\} \setminus \{q\}} t_{\mathrm{const}}^{\max}(j_u, j_q | \mathbf{z}_i)\right)$$

with

$$t_{\mathrm{const}}^{\max}(j_u, j_q | \mathbf{z}_i) = \begin{cases} tz_{iu} + \mathcal{T}_{qu}^j & \text{if } z_{i(j_u)} \neq \emptyset \wedge \mathcal{D}_{uq}^j > 0 \\ \infty & \text{otherwise.} \end{cases}$$

F. Enforcing Constraints in Phase II

The consensus phase is augmented to enforce the additional coupled constraints: unilateral dependency constraints, mutual dependency constraints, mutual exclusions, and temporal constraints.

a) Tasks with Pessimistic Bidding Strategy: During the consensus process, an agent i may find that another agent in the network outbid it for some task j_q in its bundle. When this happens, j_q as well as all tasks after j_q in bundle $\mathbf{b}_i$ are released. The tasks which are released may be depended on by other agents to satisfy the constraints associated with the tasks in their bundles. It is therefore necessary for each agent i to verify at each iteration, that the tasks in their bundle $\mathbf{b}_i$ which have a pessimistic bidding strategy have all dependency constraints satisfied. If an agent finds that it is winning a task j_q which depends on some task which is no longer assigned, it must release task j_q as well as all tasks in their bundle after j_q.

b) Tasks with Optimistic Bidding Strategy: The protocol for optimistic bidding requires that each agent keep track of the number of iterations they have been winning a task with at least one violated dependency constraint. At each iteration, agent i counts the number of satisfied constraints for each task in its bundle, and compares it to the required

number of satisfied constraints for that task. If it is less, $\nu_{i(j_q)}$ is incremented for each appropriate j_q.

Agent i also checks $\nu_{i(j_q)}$ for each task in its bundle. If $\nu_{i(j_q)} = o_{j_q}$ for any task, then j_q is released from the agent's bundle and $w_{i(j_q)}^{\mathrm{solo}}$ and $w_{i(j_q)}^{\mathrm{any}}$ are each decremented by 1. At this point, agent i has waited for o_{j_q} iterations of CBBA in hopes of acquiring all of the partnering agents that task j_q requires. The agent infers that if the constraints are not satisfied by this time, then the other agents in the network are either unable to select the tasks j_q depends on or choose not to select them because other tasks result in a greater score. Task j_q is worth nothing to agent i unless all dependency constraints are met, and so it releases that task so it can select other tasks. Agent i also releases any task appearing in its bundle after j_q, but does not decrement $w_{i(j_q)}^{\mathrm{solo}}$ or $w_{i(j_q)}^{\mathrm{any}}$ for those tasks.

c) Enforcing Mutexes: An agent i is allowed to bid on a mutexed task j_q as long as it can place a bid larger than the current winning bid of all tasks that are mutexed with j_q, and larger than the winning bid for task j_q itself. The agent placing the bid on task j_q assumes that the other agents will release the tasks mutexed with task j_q once they are informed that they were "outbid." Because of this protocol, it is possible for an agent to discover that it is currently winning a task j_u which is mutexed with another task j_q which it also believes to be assigned. Therefore it is necessary for each agent to evaluate whether they are winning a task which is mutexed with another assigned task at every iteration. Agent i is required to release task j_q from its bundle if it discovers another task j_u where $\mathcal{D}_{uq}^j = -1$ and $y_{i(j_u)} \geq y_{i(j_q)}$; it is permitted to keep task j_q at a given iteration only if

$$\left(y_{i(j_q)} > y_{i(j_u)} \vee \mathcal{D}_{uq}^j \neq -1\right) = \mathtt{true}, \ \forall u \in \mathcal{A}_j \setminus \{q\}.$$

d) Temporal Constraints: Temporal constraints may become violated during the task assignment process for a variety of reasons. Note that the task selection phase of CBBA is performed by agent i independently of all other agents, except for the information agent i posses in $\mathbf{z}_i$ and $\mathbf{y}_i$. Therefore, it is possible for several agents to bid on elements of an activity in the same bidding round, possibly resulting in conflicting arrival times. Additionally, when an agent selects the start time for a task j_q, it is only required to consider the constraints imposed by the tasks which j_q depends on. There may exist an agent k who is currently winning another task j_u which unilaterally depends on j_q. Now the start time chosen by agent i for task j_q may invalidate the start time already selected for task j_u by agent k.

Therefore, it is necessary for each agent to check the temporal constraints for all tasks in their bundle at each iteration. The method developed for enforcing timing constraints is described in this section for two cases: 1) temporal conflicts between tasks with unilateral dependency constraints, and 2) temporal conflicts between tasks with mutual dependency constraints. The process is as follows:

At each iteration, agent i checks each task $j_q \in \mathbf{b}_i$ for a

temporal constraint violation.

$$\text{tempSat}(j_q)$$
$$= \bigwedge_{u \in S_{j_q}} \left(tz_{i(j_q)} \le tz_{i(j_u)} + \mathcal{T}_{qu}^j \right) \wedge \left(tz_{i(j_u)} \le tz_{i(j_q)} + \mathcal{T}_{uq}^j \right)$$

where $S_{j_q} = \{u \in \{1, \ldots, |\mathcal{A}_j|\} : z_{i(j_u)} \ne \emptyset\}$. If $\text{tempSat}(j_q) = \texttt{true}$ the agent i keeps task j_q as it is. However, if a task j_u is found whose start time is in violation with the start time for j_q, then agent i *may* have to release task j_q.

If task j_q unilaterally depends on j_u, then agent i releases j_q, because it is violating a constraint imposed by a task assumed to be higher priority. The agent may re-bid on task j_q in the next iteration, as long as it can select a start time that does satisfy the constraints. If task j_u unilaterally depends on j_q, then agent i keeps task j_q in its bundle. The agent assumes that the agent winning task j_u will also realize that j_q and j_u are conflicted, and release task j_u.

If task j_q and j_u are mutually dependent, then a special procedure is followed: Assume that the score for a given task is monotonically decreasing within its time window. Then the agent arriving earliest (with respect to the start time of the task) is required to release that task. In this instance, it is assumed that the agent arriving later in their task's time window would have chosen an earlier start time if it were possible, since the score decays with time. Given that agent i is winning task j_q and agent k is winning task j_u, and the start times are conflicted, then agent i releases task j_q if

$$tz_{i(j_q)} - \tau_{j_q}^{\text{start}} \le tz_{i(j_u)} - \tau_{j_u}^{\text{start}}$$

Now each of the cases for the (j_q, j_u) task pair have been accounted for when there is a temporal constraint violation between them. If an agent is required to release a task j_q because of a temporal constraint violation, and task j_q has an optimistic bidding strategy, then $w_{i(j_q)}^{\text{solo}}$ and $w_{i(j_q)}^{\text{any}}$ are each decremented by one.

V. NUMERICAL RESULTS

A. Simulation Setup

Coupled-constraint CBBA is compared to the baseline CBBA using a simulation to determine the merit of the assignments achieved. The mission scenario that is simulated is a cooperative search, track, and engage mission. Participating in the mission, are three types of agents. Agents of type A are Neutralizing Uninhabited Aerials Vehicles (NUAVs) which are capable of engaging hostile targets. Agents of type B are Sensor UAVs (SUAVs), and they are capable of providing the NUAVs with accurate measurements during an engagement. Finally, agents of type C are SUAVs, but are only capable of image capture; they cannot provide measurements to a NUAV. Agent capabilities are described in the Table II.

There exist several objects of interest in the mission environment. Some of the objects are confirmed hostiles, and it is desired that all hostiles be engaged. Some of the objects are of unconfirmed identity, and more intelligence is desired

TABLE II
AGENT PARAMETERS FOR SIMULATED MISSION SCENARIO

Agent Parameters		
	Units Available	3
Type A: NUAV	Capability	Neutralization
	Max Velocity	50 m/s
	Units Available	5
Type B: SUAV 1	Capability	Sensing for NUAV
	Max Velocity	25 m/s
	Units Available	8
Type C: SUAV 2	Capability	Sensing- Image only
	Max Velocity	15 m/s

regarding these objects. The two different types of objects each have a specific activity associated with them.

Activity type 1 is a *service hostile* activity: hostile targets may be engaged, which requires one agent of type A to strike the target and one agent of type B to provide position measurements. The strike portion of the activity takes 120 seconds, and the NUAV and the SUAV must plan to arrive within 20 seconds of each other. If one type A and one type B agent are assigned to a service hostile activity, a type C agent may also be assigned to perform a damage assessment (DA). The DA takes 180 seconds, and must begin at least 60 seconds after the strike is completed. If there is no strike package available to engage the hostile target, a SUAV may be sent to observe the hostile, but it is assumed that the reward for performing the intelligence gathering task is much less than that for performing the strike task. The maximum score for a strike task is 100, the maximum score for a DA is 50, and the maximum score for a intelligence gathering task is 10. The time window for the activity is 600 seconds.

Activity type 2 is a tracking activity and involves visiting one of the objects with unknown identity. This type of activity contains a single element which is an information gathering task that can be performed by either a type B agent or a type C agent. The duration is 60 seconds, and the maximum score for a tracking activity is 10. The dependency and temporal constraint matrices for an activity of type 1 and type 2 are

$$\mathcal{D}_1 = \begin{bmatrix} 0 & 1 & 1 & -1 \\ 1 & 0 & 1 & -1 \\ 0 & 0 & 0 & -1 \\ -1 & -1 & -1 & 0 \end{bmatrix} \quad \mathcal{T}_1 = \begin{bmatrix} 0 & 20 & -180 & \infty \\ 20 & 0 & -180 & \infty \\ \infty & \infty & 0 & \infty \\ \infty & \infty & \infty & 0 \end{bmatrix}$$
$$\mathcal{D}_2 = \begin{bmatrix} 0 \end{bmatrix}, \qquad \mathcal{T}_2 = \begin{bmatrix} 0 \end{bmatrix}$$

The simulation environment is 10×10 km. Activity locations and agent initial positions are generated randomly, uniform in the environment. The number of activities is varied as a parameter, with half of the activities being of type 1, and half type 2. The time windows for each activity are chosen such that the start time is uniformly random on [0, 300] seconds. A maximum bundle length, $L_t = 4$ is used.

The scoring for a given assignment is calculated such that only tasks which are assigned without constraint violation count toward the assignment score. Also, the baseline CBBA has no way to account for timing constraints, so the only way to enforce them is to shrink the time window for the two elements of the neutralization task to be 20 seconds long, and set the time window of the DA to begin after the

Fig. 2. Coupled-constraint CBBA vs. Baseline CBBA, Assignment Score

Fig. 4. Coupled-constraint CBBA vs. Baseline CBBA, Computation Complexity

Fig. 3. Coupled-constraint CBBA vs. Baseline CBBA, Number of Tasks Assigned
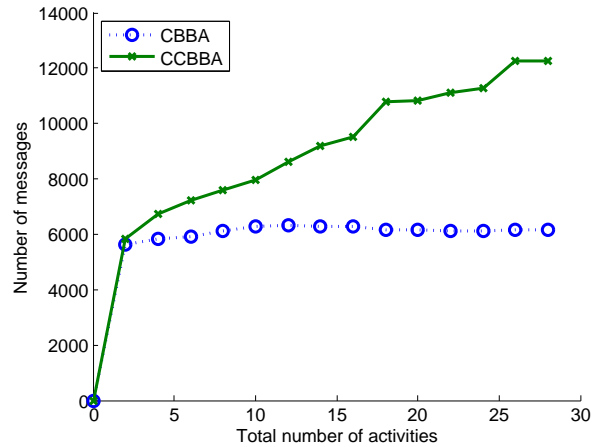
Fig. 5. Coupled-constraint CBBA vs. Baseline CBBA, Comm. Complexity

neutralization window ends.

B. Results

The number of activities is varied as a parameter and 80 trials are simulated for each case. Fig. 2 compares the average assignment scores of the baseline CBBA and CCBBA. Notice that the score for CBBA flattens out as the number of activities increases beyond 10. The reason for this behavior is: As N_t/N_u grows larger, each agent has increasingly more task options available. Given that an agent selects a task, j_q, the probability that all tasks which j_q depends on are also assigned decreases as N_t/N_u increases when using the baseline CBBA. However, notice that the score for CCBBA-generated assignments increases with the number of tasks at a much higher rate compared to CBBA, and does not posses this tendency to flatten out suddenly. This is because CCBBA is explicitly enforcing the coupled constraints so that agents select tasks that are actually valuable.

Fig. 3 further validates the advantage of using CCBBA for large N_t/N_u cases. The figure compares the total number of tasks assigned by each algorithm as well as the number of feasible tasks assigned. A feasible task is defined to be a task which satisfies all constraints in the problem. Notice that the total number of tasks assigned is roughly the same for each approach. However, for the baseline CBBA, the percentage of tasks assigned which satisfy the constraints decreases as

the number of activities increases, thus adversely affecting the assignment score for the baseline CBBA. The figure also shows that for CCBBA, all assigned tasks satisfy the constraints, which is reflected in the assignment score.

Fig. 4 shows the increase in computation for CCBBA compared to CBBA. The metric for comparing computation is the required number of score calculations to arrive at an assignment summed over the whole network. The comparison of computation shows that for this mission scenario, the computation requirements are approximately linear for CBBA, and approximately quadratic for CCBBA.

Fig. 5 compares the communication requirements for the two algorithms. The metric for comparing communication is the total number of messages parsed during the assignment process summed over the network. This figure shows that for this choice of mission parameters, CBBA has approximately constant communication requirements with respect to the number of tasks, whereas CCBBA has approximately linear communication requirements. The additional complexity is associated with the optimistic bidding process, because agents must put forth additional effort to discover that a task will not have the required number of satisfied constraints.

CBBA is known to have polynomial computation and communication requirements. CCBBA is certainly more computationally expensive compared to CBBA, and requires

additional message passing. However, CCBBA still provides a tractable framework for complex missions involving large teams and many tasks. In this particular mission scenario, there were 16 agents and up to 30 activities, (75 tasks with coupling constraints), and the assignments were generated in less than 40 seconds (in the Matlab environment with all agents' computations being performed by a single CPU).

VI. CONCLUSIONS

Coupled-constraint CBBA (CCBBA) was developed as an efficient mechanism for solving the decentralized constrained task allocation problem. These algorithmic extensions to CBBA expand the types of constraints that CBBA can characterize to include unilateral dependency constraints, mutual dependency constraints, mutual exclusion, and temporal constraints. Numerical examples validate that CCBBA enables incorporation of a wide range of mission constraints, providing measurable performance improvement to the baseline CBBA. Future work will include integration of the presented CCBBA into the Asynchronous CBBA framework [15] that has been developed for the most efficient implementation of CBBA in terms of computation and communication.

ACKNOWLEDGMENTS

REFERENCES

[1] J. Bellingham, M. Tillerson, A. Richards, and J. How, "Multi-Task Allocation and Path Planning for Cooperating UAVs," in *Proceedings of Conference of Cooperative Control and Optimization*, Nov. 2001.
[2] C. Schumacher, P. Chandler, and S. Rasmussen, "Task allocation for wide area search munitions," in *Proceedings of the American Control Conference*, 2002.
[3] A. Casal, "Reconfiguration planning for modular self-reconfigurable robots," Ph.D. dissertation, Stanford University, Stanford, CA, 2002.
[4] Y. Jin, A. Minai, and M. Polycarpou, "Cooperative Real-Time Search and Task Allocation in UAV Teams," in *IEEE Conference on Decision and Control*, 2003.
[5] D. Turra, L. Pollini, and M. Innocenti, "Fast unmanned vehicles task allocation with moving targets," in *Proceedings of the IEEE Conference on Decision and Control*, Dec 2004.
[6] M. Alighanbari, "Task assignment algorithms for teams of UAVs in dynamic environments," Master's thesis, Massachusetts Institute of Technology, 2004.
[7] T. W. McLain and R. W. Beard, "Coordination variables, coordination functions, and cooperating-timing missions," *Journal of Guidance, Control, and Dynamics*, vol. 28(1), pp. 150–161, 2005.
[8] D. A. Castanon and C. Wu, "Distributed algorithms for dynamic reassignment," *IEEE Conference on Decision and Control*, vol. 1, pp. 13–18 Vol.1, 9-12 Dec. 2003. [Online]. Available: 10.1109/CDC.2003.1272528
[9] J. Curtis and R. Murphey, "Simultaneous area search and task assignment for a team of cooperative agents," in *Proceedings of AIAA Guidance, Navigation, and Control Conference and Exhibit*, 2003.
[10] T. Shima, S. J. Rasmussen, and P. Chandler, "Uav team decision and control using efficient collaborative estimation," in *Proceedings of the American Control Conference*, 8-10 June 2005, pp. 4107–4112 vol. 6. [Online]. Available: 10.1109/ACC.2005.1470621
[11] M. Alighanbari and J. How, "Robust and Decentralized Task Assignment Algorithms for UAVs," Ph.D. dissertation, MIT, 2007.
[12] H.-L. Choi, L. Brunet, and J. P. How, "Consensus-based decentralized auctions for robust task allocation," *IEEE Trans. on Robotics*, vol. 25 (4), pp. 912 – 926, 2009.
[13] S. Ponda, J. Redding, H.-L. Choi, J. P. How, M. A. Vavrina, and J. Vian, "Decentralized planning for complex missions with dynamic communication constraints," in *American Control Conference (ACC)*, July 2010.
[14] H.-L. Choi, A. K. Whitten, and J. P. How, "Decentralized task allocation for heterogeneous teams with cooperation constraints," in *American Control Conference (ACC)*, July 2010, pp. 3057–3062.
[15] L. B. Johnson, S. Ponda, H.-L. Choi, and J. P. How, "Improving the efficiency of a decentralized tasking algorithm for UAV teams with asynchronous communications," in *AIAA Guidance, Navigation, and Control Conference (GNC)*, August 2010.

APPENDIX

A. Examples of Encoding Real-World Constraints

Encoding the constraints into the dependency matrix and the temporal constraint matrix may not be intuitive. This section provides a couple of different examples of mission scenarios, and describes the corresponding constraints matrices. Details on more examples can be found in [**?**, Section 3.9].

a) Required Cooperation: Many real-world scenarios involve a set of tasks, all of which must be assigned for any of them to be valid, and they must have simultaneous arrival. The dependency matrix for a set of mutually required tasks of arbitrary size are given by $\mathcal{D} = \mathbf{1}_{N_t \times N_t} - I_{N_t}$, and temporal constraint matrix in this case is $\mathcal{T} = \mathbf{0}_{N_t \times N_t}$.

b) Not During: Consider a scientific operation where a series of measurements are taken by a fleet of robots. Suppose there are two measurements that are to be taken: measurement task A, and measurement task B. Both measurements must be taken for either of them to be valid, however, they cannot be taken simultaneously, due to equipment interference. The constraints can be encoded by splitting one of the tasks into two instances - consider B into B^- and B^+ for example. Then, task B^- depends on task A, task B^+ depends on task A, and A depends on either B^- or B^+, which give

$$\mathcal{D} = \begin{bmatrix} 0 & 1 & 1 \\ 2 & 0 & -1 \\ 2 & -1 & 0 \end{bmatrix}, \mathcal{T} = \begin{bmatrix} 0 & \infty & -d_A \\ -d_B & 0 & \infty \\ \infty & \infty & 0 \end{bmatrix}.$$

c) Super-Additive Score Structure: Sometimes, it is desirable to send two UAVs to track a target if possible, severely reducing the probability of being evaded. This situation often leads to a super-additive score structure for which the combined score is greater than the sum of the individuals. To encode this score structure, combined assignment is represented by two tasks, each worth half of the combined score. A third task representing single assignment is introduced; this task is to be mutexed with each of the two tasks representing combined assignment, which give constraint matrices

$$\mathcal{D} = \begin{bmatrix} 0 & 1 & -1 \\ 1 & 0 & -1 \\ -1 & -1 & 0 \end{bmatrix}, \mathcal{T} = \begin{bmatrix} 0 & 0 & \infty \\ 0 & 0 & \infty \\ \infty & \infty & 0 \end{bmatrix}.$$