Constrain, Correspond, Correct: Distributed Game-Theoretic Data Association for Assignment Games on Multimodal Sensing Grids

Robert J. Ravier
Sarcos Group, LC
Salt Lake City, UT, USA
r.ravier@palladyneai.com

Denis Garagić
Sarcos Group, LC
Salt Lake City, UT, USA
d.garagic@palladyneai.com

Travis Galoppo
Sarcos Group, LC
Salt Lake City, UT, USA
t.galoppo@palladyneai.com

Rex Jameson
Sarcos Group, LC
Salt Lake City, UT, USA
r.jameson@palladyneai.com

Bradley J. Rhodes
Crystalytyx LLC
Burlington, MA, USA
brad@crystalytyx.com

Peter Zulch
AFRL/RI
Rome, NY, USA
peter.zulch@us.af.mil

***Abstract*—Assignment games are a promising framework for autonomous management of a multimodal Tactical Sensing Grid (TSG). They provide theoretical guarantees and exhibit excellent empirical performance in maintaining custody of all observed targets in a scene. However, TSG sub-grid initiation of an assignment game requires multimodal data association between playing nodes (the sub-grid). Playing nodes must achieve consensus on target labels and identities before game play proceeds. Using the locally centralized communication network assumed by an assignment game, we propose the Triple-C distributed method for solving this association problem. The method we propose is suitable for fairly general scenarios, requiring that each node have an intrinsic notion of what constitutes an outlier and what constitutes similarity. At the core of the Triple-C method is a collection of parallelizable assignment games played between two nodes. Triple-C yields theoretical guarantees on output association consistency. We evaluate the performance of the Triple-C method on multiple simulations, showing that it computes associations quickly and accurately, thus enabling a TSG to maintain full situational awareness in a two-vehicle scenario.**

***Keywords*—data association, game theory, reinforcement learning, multimodal, situational awareness, distributed algorithms**

I. Introduction

The problems of constructing methods that enable a multi-platform tactical sensor grid (TSG), where each platform observes one or multiple modalities of sensor information (e.g., images, radar, passive RF), to autonomously maintain and maximize situational awareness (SA) over a possibly time-varying area of interest remains of immense interest and challenge. This inevitably requires multiscale solutions; each sensor platform must be able to manipulate its sensors as needed to accurately detect, track, classify, and identify objects of interest, and the platforms themselves must communicate with one another to properly allocate their respective resources to ensure the TSG as a whole achieves SA.

Our prior work to this end has shown that the combination of reinforcement learning (RL) [1, 2, 3], sensor fusion [4, 5, 6, 3], and game theory [7] yields an effective solution capable of maintaining track custody of multiple objects in an area of interest. In particular, in [2, 3], we showed that online reinforcement learning is capable of quickly learning policies for manipulating individual sensors of multiple modalities for maintaining SA, and that fusing multimodal sensor data improves the ability for said policies to generalize. We used game theory in [7] to combine policies for individual sensor platforms in a TSG via *assignment games*, which assign to each sensor a subset of the observed objects in a scene in situations in which the individual platforms cannot possibly maintain custody of all objects for which they are aware. We showed that there exists a broad class of utilities that result in *potential* games, in which a number of algorithms are guaranteed to converge to a Nash equilibrium. This guarantee ensures that all objects in an assignment game are fully allocated to the collective of platforms participating in the game, i.e. there is no risk for known objects to be lost. We then demonstrated that our assignment games enable a simulated four node TSG to maintain full custody of multiple vehicles moving in opposing directions, an impossible task for any individual platform.

To apply our assignment game methodology in practice, we propose a solution for the cross-platform data association problem that went unaddressed in [7]. It is perhaps easiest to think of this in terms of images. Suppose multiple cameras observing roughly the same area each observe their own collection of objects. In this setting, the cross-platform data association problem seeks to determine, for a given subset of cameras, which objects are in the field of view of each camera, and for each such object, determine the correspondence between its label in one camera (if assigned an identification number) with those in the other cameras. Such accurate correspondence information is needed to play an assignment game; if such labels are inconsistent between platforms, the guarantee of known objects being lost during an assignment game is lost. For example, if an assignment game seeks to allocate two objects to two platforms, but each platform has opposite labels for the objects, the assignment game may yield that both platforms track the same object, leaving the other to be lost.

Though the data association problem has been well-studied (see Section II for a brief overview of possible directions), both the context of [7] as well as our interest in developing methods that perform well in volatile, uncertain, complex, and ambiguous

This paper was funded under AFRL contract FA8750-22-C-1005.

environments presents unique challenges and constraints. To ensure that a given TSG maintains SA, any assignment game must be completed as quickly as possible so as to guarantee that the individual sensor platforms can manipulate their sensors and not lose any observed objects; this means that any prerequisite data association method must be similarly fast. Any such method should also be as resource non-intensive as possible so as to allow for smooth operation in tandem with computationally intense algorithms (e.g., large neural networks) on SWaP constrained hardware that may only be able to communicate a small subset of relevant information. To that end, optimal methods must leverage distributed algorithms to properly offset said hardware constraints. Prior work on achieving SA in compute-limited scenarios (e.g., drones) has fruitfully addressed these issues in several applications [8, 9, 10].

We provide a solution to our cross-platform data association problem, which we dub the Triple-C method, based on an adaptation of our previously proposed assignment game methodology from [7]. At its core, Triple-C leverages a collection of assignment games that include a class of utilities for which Nash equilibria ensure assignment of a corresponding label to all observed objects. We make explicit use of a game network topology assumption that is always valid within the context of our data-association problem, irrespective of the actual TSG-wide topology. This allows us to significantly reduce the number of computations required while simultaneously allowing us to distribute the computations over all game playing nodes.

The paper proceeds as follows. We review prior work in Section II, and give a rigorous overview of the problem and related context in Section III. We detail our proposed algorithm in Section IV. In Section V, we discuss the simulation environment we use, and detail our resulting experiments in Section VI. We make concluding remarks in Section VII.

II. Prior Work on Data Association

The data association problem is a subject with a vast literature well deserving of a review paper on its own. The problem of data association naturally occurs within the tracking problem [11, 12, 13, 14] as observations at a given time need to be associated to those made at prior times. There are many different techniques that can be used based on the structure of the problem of interest. Given limited space, we briefly review some methods relevant to our multimode scenario.

Unsurprisingly, data association techniques for multi-platform methods have heavily focused on camera networks [15, 16, 17]. Provided sufficient knowledge of camera parameters, such as field of view, this allows for the construction of association algorithms based on geometric techniques [18], though our Triple-C approach does not assume that such information is always available. Recent work has proposed solutions based on simultaneous tracking of appearance features and estimation of camera network topology [19, 20]. Our approach obviates the need for sensor network topology estimation and reduces overall computational burden. Other work has proposed non game-theoretic optimization techniques [21, 22]. The literature on camera network association problems is itself large; we defer the reader to [14] for an extensive review on recent developments.

III. Problem Overview

A. Game Theory and Abstract Assignment Games

Recall that a game with n players, in the game-theoretic sense, is a collection of 2-tuples (A_i, U_i) for $1 \leq i \leq n$, where A_i denotes the *action set* of player i and $U_i\colon \prod_{i=1}^{n} A_i \to \mathbb{R}$ is player i's *utility function*. Here, $\prod_{i=1}^{n} A_i$ denotes the Cartesian product of the sets $A_1, \dots, A_n$, though usually the order of arguments in utility functions are ignored. In the following we assume that the action sets are all finite. For each $t \in \mathbb{Z}_+$, each player selects an action $a_i^t \in A_i$, making each player's utility $U_i(a_1^t, \dots, a_n^t)$. In game theory, the goal is often to compute *Nash equilibria* of a given game, the analogue of local maxima for multiple objective functions; letting a_{-i} denote an arbitrary element of $A_1 \times \dots \times A_{i-1} \times A_{i+1} \times \dots \times A_n$, we say that an element $a^* \in \prod_{i=1}^{n} A_i$ is a Nash equilibrium of the game $\{(A_i, U_i)\}$ if we have

$$U_i(a_i^*, a_{-i}^*) \geq U_i(a, a_{-i}^*) \text{ for all } a \in A_i, \text{ for } 1 \leq i \leq n.$$

We now define two different types of *abstract assignment games*. Given two nonempty finite sets $X = \{x_1, \dots, x_{k_1}\}$ and $Y = \{y_1, \dots, y_{k_2}\}$, the elementwise assignment game has as players the elements of X, the actions of each x_i are assignments to individual elements of Y (so the collective action of all players at time t can be thought of as a function $g_t\colon X \to Y$) and the utilities U_i are of the form

$$U_i(g_t) = -f_i\big(x_i, g_t(x_i)\big) - C\,\frac{C^{k_2 - \mathrm{card}\left(\cup_{i=1}^{k_1}\{g_t(x_i)\}\right)} - 1}{C-1} \quad (1)$$

for a constant $C > 1$ and functions $f_i\colon X \times Y \to \mathbb{R}_+$. Similarly, the subset assignment game has as players the elements of X, the actions of each x_i are assignments to subsets of Y (so the collective action of all players at time t can be thought of as a function $h_t\colon X \to 2^Y$ (here, 2^Y denotes the power set of Y)) and the utilities U_i' are of the form

$$U_i'(h_t) = -f_i'\big(x_i, h_t(x_i)\big) - C\,\frac{C^{k_2 - \mathrm{card}\left(\cup_{i=1}^{k_1} h_t(x_i)\right)} - 1}{C-1} \quad (2)$$

for a constant $C > 1$ and functions $f_i'\colon X \times 2^Y \to \mathbb{R}_+$. The right most terms of both Equations 1 and 2 act as a joint cardinality penalty for all players; all utilities increase provided that the assignments cover more of Y.

We showed in [7] that both types of assignment games have a number of desirable properties. Both the elementwise and subset assignment games are *potential games*, for which multiple algorithms guarantee convergence to Nash equilibria [23, 24]. In this work, we focus on *spatial adaptive play* (SAP), an algorithm similar in spirit to the simulated annealing optimization algorithm [25]. In SAP, the game initializes with each player selecting their actions uniformly at random. At each time t, only one player is chosen at random to select a new action; other players keep the same action. With a scalar $\beta > 0$, assuming that player i is selected to update its action, it samples an action from the Gibbs measure

$$p_i^t(a_i) \propto exp\Big(\beta U_i(a_i, a_{-i}^{t-1})\Big). \quad (3)$$

In practice, β is a (generally increasing) function of t, where lower values result in sampling from distributions closer to uniform and higher values result in sampling from distributions that concentrate around $a \in A_i$ that maximize U_i.

For both types of assignment games, there exist a broad class of utilities for which Nash equilibria give assignments of maximum cardinality. Specifically, if the f_i' in the subset assignment game have magnitudes uniformly bounded by some constant M, then by selecting $C \geq M$, any Nash equilibrium h^* satisfies $\cup_{i=1}^{k_1} h^*(x_i) = Y$, i.e. the assignments cover all of Y. Similarly, if the f_i in the elementwise assignment game have magnitudes uniformly bounded by some constant M, then by selecting $C \geq M$, then for any Nash equilibrium g^*, the set $\cup_{i=1}^{k_1} \{g^*(x_i)\}$ has cardinality $\min\big(\mathrm{card}(X), \mathrm{card}(\mathrm{Y})\big)$, i.e. the image of g^* is of maximum cardinality.

The subset assignment game has an additional property of minimizing overlap of resulting assignments. If the f_i' are also monotone ($f_i'(x, A) < f_i'(x, B)$ if A is a strict subset of B) in addition to being uniformly bounded, then at a Nash equilibrium h^*, we have $h^*(x_i) \cap h^*\big(x_j\big) = \emptyset$ for $i \neq j$.

B. Practical Assignment Games for Tactical Sensor Grids

Both the potential game and cardinality properties of each type of assignment games result in effective and practical solutions for the following sensor-to-target assignment problem for a TSG. Consider a situation where n sensors that are mutually aware of k objects in their covered area of interest wish to allocate all k objects to individual sensors for further tracking. By the previous section, the sensors can play a subset assignment game that will yield a solution that will allocate each object to at least one sensor for further monitoring. Objects can also be fully allocated to sensors by playing an elementwise assignment game, though the specific game played depends on n and k. For $n \geq k$, a game assigning individual sensors to individual objects will work; otherwise, a game assigning individual objects to individual sensors will do. Since all games played are potential games, algorithms like spatial adaptive play will theoretically converge to such an assignment, which we showed in [7] to occur relatively quickly in practice. It thus remains to discuss when to play an assignment game.

We model a TSG as a (possibly time-varying) graph, where the vertices of the graph correspond to individual sensor platforms in the TSG, and the existence of an edge between any two vertices means that it is possible for the two corresponding sensor platforms to communicate with one another. In determining whether and when to play an assignment game, we assume that each node has both its own *request game* criteria and *accept game* criteria. The request game criteria are a variable condition in which a sensor platform decides to propose an assignment game with other platforms due to its inherent inability to continue tracking all of the observed objects it perceives. For example, a gimbal camera platform's request game criteria may consist of observing a pair of objects moving in opposite directions, for which the gimbal camera would inevitably have to choose between which object it continues to track.

The accept game criteria is a condition for which, upon receiving a request game message, a given sensor platform decides whether it can deviate from its current actions without jeopardizing its own interests. For example, a platform's accept game criteria may be determining if its actions can change in a way without compromising its ability to be aware of an object that it considers of high value to track, or whether it is currently engaged in an assignment game.

With these two criteria in mind, a subset of nodes initializes an assignment game in a *locally* centralized fashion dependent on the proximity of sensor platforms to one another and the ability for platforms in close proximity to communicate with one another. Specifically, if an individual sensor platform's request game criteria is met, that platform transmits a message to other platforms in its vicinity that it can communicate with, requesting to play an assignment game. The messaged nodes evaluate their own accept game criteria, upon which they communicate their acceptance back to the initial node that requested the game. The requesting node compiles the list of accepting nodes, and transmits the necessary information for itself and the accepting nodes to play an assignment game. Fig. 1 illustrates this game play procedure. Of particular note is the resulting topology of the subgraph of nodes playing the game; the subgraph is a tree with star-shaped topology. The root is the node that requested the game, and the leaves are the nodes that accepted the game. We will make explicit use of this fact in Section IV.

C. Problem Statement

Fig. 1 depicts the assignment game under consideration in this paper. Of the N TSG nodes, n are playing this game. Node r is the requesting node that initiates this game (where $1 \leq r \leq n$. The n playing nodes are collectively aware of T targets in the scene, with node i being aware of T_i targets (enumerated by $t_1^i, \ldots, t_{T_i}^i$). We do not assume that each node observes a disjoint

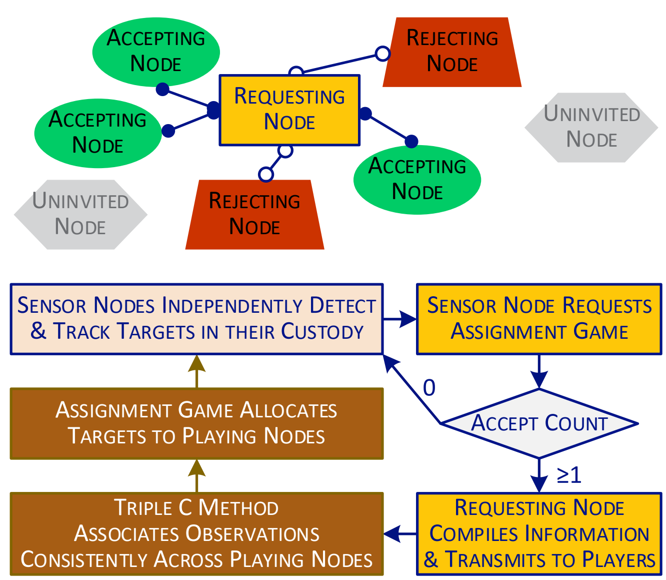

Fig. 1. Example of the assignment game playing sub-grid topology (*Top*) and flow (*Bottom*).

set of targets, so the sum of the T_i is at least T. This assignment game thus aims to assign targets $t_1^r, \dots, t_{T_r}^r$ to the n nodes (including node r).

As defined above, this assignment game is not immediately playable. The critical missing element is an unambiguous effective association between targets observed by each playing node, because the abstract assignment game inherently assumes existence of such. That is, an abstract assignment game assumes that a given target t has a known label common across all nodes. Therefore, for valid assignment game play, we must solve the following problem: establish, for each accepting node (a), an association between the targets a observes and the targets r observes. More formally: for $a \neq r$, how do $t_1^a, \dots, t_{T_a}^a$ correspond to $t_1^r, \dots, t_{T_r}^r$, if at all? To be effective, associated targets must correspond to the same physical object in the scene. With effective established associations, valid game play proceeds using r's target labels.

IV. The Triple-C Method: Constrain, Correspond, Correct

Our proposed inter-node target association solution exploits the graph structure depicted in Fig. 1 to avoid having to compute unambiguous associations globally. Note that, without assumptions about the time-varying topology of the TSG graph as a whole, the initialization of an assignment game leads to a *star-shaped* graph topology, with the requesting node central. This topology allows a simpler approach for establishing the needed associations by computing pairwise associations between the (central) requesting node and each accepting node. The problem is now one of solving a smaller association problem for each accepting node. Modern parallel compute capabilities are adept at this form of problem.

In the following, we outline some mild assumptions regarding the data and methods available to each node. Then we detail our general game theoretic solution, making use of the abstract assignment game framework reviewed in Section III.

A. Assumptions

We assume that sensor node i can observe M_i modalities $m_1^i, \dots, m_{M_i}^i$. We represent an observed target t_j^i ($1 \leq j \leq T_i$) as a collection of summary statistics in each modality:

$$t_j^i = \bigcup_{k=1}^{M_i} S_k^{i,j} = \bigcup_{k=1}^{M_i} \bigcup_{\ell=1}^{L_{M_i}} \{s_{k,\ell}^{i,j}\}. \tag{4}$$

Here $S_k^{i,j}$ is the set of summary statistics of target t_j^i in the kth modality.

Each node possesses additional utility functions conditioned on one or all of the known targets. One is an outlier utility $O_i(\cdot \,|t_1^i, \dots, t_{T_i}^i)$, and the others comprise a family of dissimilarity utilities $P_i(\cdot, t_j^i)$, one for each t_j^i. Without loss of generality, we assume that $0 \leq O_i, P_i \leq 1$. Intuitively, higher values of $O_i(x|t_1^i, \dots, t_{T_i}^i)$ suggest that x is not similar to any of the observed data. Likewise, lower values of $P_i(x, t_j^i)$ suggest that x is closer to t_j^i.

We also assume that, for requesting node r and any accepting node a, there are functions $f_{a,r}$ and $f_{r,a}$. $f_{a,r}(t_k^r)$ is a transformation mapping the summary statistics of t_k^r into a collection of summary statistics of the same form as those of t_j^a. Similarly for $f_{r,a}(t_j^a)$, mapping from t_j^a to t_k^r. Note that neither $f_{a,r}$ nor $f_{r,a}$ establish any association; they merely map modalities to modalities between TSG nodes. This can be rather broad in practice. However, for two EO camera-bearing nodes (as presented in Section V), the identity will suffice. We do not assume that $f_{a,r}$ and $f_{r,a}$ are inverses or even pseudoinverses.

B. Constrain the Targets in Play

For requesting node r and any accepting node $a \neq r$, it is not clear *a priori* that node a is currently observing any of the t_k^r. To initiate the association game, requesting node r transmits all of its known targets t_k^r to accepting node a. Node a transforms these t_k^r into corresponding collections of statistics $f_{a,r}(t_k^r)$ that are directly comparable to existing targets t_j^a of node a. Next, node a evaluates $f_{a,r}(t_k^r)$ against the t_j^a via $O_i(f_{a,r}(t_k^r)|t_1^i, \dots, t_{T_i}^i)$. If any O_i value is below some platform specific threshold, ε_a, then node a concludes that it is not observing that target, t_k^r. After full evaluation, node a sends the list of t_k^r it is observing back to node r. This list provides an initial constraint on the targets mutually perceived by node r and node a. These targets are thus eligible assignees in an assignment game between node r and node a.

Note that the reverse procedure can also occur, pending resource and time constraints. Namely, node r processes and returns a list of the t_j^a that it deems relevant to the assignment game. We denote the constrained lists of targets that are relevant for future game play by $\hat{t}_k^r$ and $\hat{t}_j^a$. Note that if one of the lists is empty, then node a no longer participates in the assignment game since it does not observe any targets of interest to node r.

C. Establishing Initial Correspondence Between Targets

With the constrained target lists, the nodes now play an elementwise assignment game involving the sets $\{\hat{t}_k^r\}_{k=1}^K$ and $\{\hat{t}_j^a\}_{j=1}^J$. Comparison between set cardinalities determines which game to play. Following the notation of Equation 1, the utilities of the assignment game assigning $\hat{t}_k^r$ to $\hat{t}_j^a$ are:

$$U_k^r(g) = -P_a\left(f_{a,r}(\hat{t}_k^r), g(\hat{t}_k^r)\right) - \ell(C, J, K, \{\hat{t}_k^r\}_{k=1}^K, g), \tag{5}$$

while those for assigning $\hat{t}_j^a$ to $\hat{t}_k^r$ are:

$$U_j^a(g') = -P_r\left(f_{r,a}(\hat{t}_j^a), g'^{(\hat{t}_j^a)}\right) - \ell\left(C, K, J, \{\hat{t}_j^a\}_{j=1}^J, g'\right). \tag{6}$$

Here, the function $\ell(C, J, K, \{\hat{t}_k^r\}_{k=1}^K, g)$ is the cardinality penalty, $C^{\frac{C^{J-\mathrm{card}(\cup_{i=1}^{K}\{g(\hat{t}_i^r)\})}-1}{C-1}}$, for mapping $\hat{t}_k^r$ to $\hat{t}_j^a$; $\ell\left(C, K, J, \{\hat{t}_j^a\}_{j=1}^J, g'\right)$ is similar. Intuitively, higher values of $U_k^r(g)$ mean that the image of g has higher cardinality, and that $\hat{t}_k^r$'s transformed summary statistics under $f_{a,r}$ are similar to those of its corresponding target $g(\hat{t}_k^r)$. (Note that meaningful comparison requires the $f_{a,r}$. In an abstract sense, the $f_{a,r}$ are parallel transports [26]). The utilities $U_j^a(g')$ behave similarly.

The relationship between J and K uniquely determines the game to play. If $J < K$, then game play assigns $\hat{t}_j^a$ to $\hat{t}_k^r$. If $K < J$, game play assigns $\hat{t}_k^r$ to $\hat{t}_j^a$. If $J = K$, nodes play either game. The following theoretical result motivates this choice:

Theorem 1: $U_k^r(g)$ and $U_j^a(g')$ define potential games. Furthermore, if g_* and g'_* denote Nash equilibria of the $U_k^r(g)$ and $U_j^a(g')$ games respectively, then the cardinality of their images $\cup_{i=1}^{K} \{g_*(\hat{t}_i^r)\}$, $\cup_{i=1}^{J} \{g'_*(\hat{t}_i^a)\}$ equals $\min(J, K)$.

Theorem 1 is a special case of the results presented in [7]. The cardinality property inspires game choice because the Nash equilibrium assignments are guaranteed to be injections. Per standard properties of injections, this result induces a bijection from a non-empty subset of $\hat{t}_k^r$ to a non-empty subset of $\hat{t}_j^a$, which we denote by $g_{a,r}$ (and $g_{r,a}$ denotes the inverse).

Thus, we now have a collection of non-trivial bijections $g_{a,r}$, one for each accepting node a. This provides sufficient information to play the sensor to target assignment game by propagating the labels of t_k^r via $g_{a,r}$. Through the bijections, node a knows unambiguously whether a given t_k^r corresponds with one of its t_j^a. Furthermore, this reduces the resulting search space for each accepting node, potentially dramatically. An accepting node needs only focus on those t_k^r for which it perceives corresponding targets.

D. Correct Correspondences via Anomalies & Consistency

Computation of bijections $g_{a,r}$, $g_{r,a}$ uses only pairwise information shared between requesting node r and accepting node a. This omits use of information available to r from other as. A final algorithmic step uses the dissimilarity utilities to correct the bijections via thresholds specific to node r in two procedures. First, *anomaly correction* computes $P_r\left(f_{r,a}\left(g_{a,r}(t_k^r)\right), t_k^r\right)$ for each t_k^r in the domain of $g_{a,r}$. Results above some platform specific threshold δ_r^o cause pruning of $g_{a,r}$ (and by proxy, $g_{r,a}$) to remove any reference to t_k^r. This removes associations that node r views as invalid because they are sufficiently dissimilar.

Second, *consistency correction* compares, for each t_k^r, all associated t_j^a from the computed bijections. If there is at most one such t_j^a, nothing is done. Otherwise, we keep two lists: one of accepting nodes L_A, and one of unordered pairs of accepting nodes $L_{A\times A}$. For each unordered pair $\hat{a}, \hat{a}' \in \hat{A} \subseteq A$ of distinct accepting nodes for which t_k^r lies in the domain of both $g_{\hat{a},r}$ and $g_{\hat{a}',r}$, we compute dissimilarities

$$P_r\left(f_{r,\hat{a}}\left(g_{\hat{a},r}(t_k^r)\right), f_{r,\hat{a}'}\left(g_{\hat{a}',r}(t_k^r)\right)\right).$$

If these dissimilarities are above some platform specific δ_r^c, then we append the pair $(\hat{a}, \hat{a}')$ to $L_{A\times A}$, and add both $\hat{a}$ and $\hat{a}'$ to L_A if L_A does not already contain them (duplicates are not allowed). After iterating through all unordered pairs, we sort L_A in decreasing order based on the previously computed values of $P_r\left(f_{r,\hat{a}}\left(g_{\hat{a},r}(t_k^r)\right), t_k^r\right)$. We then iterate through the $\hat{a} \in L_A$, pruning the $g_{\hat{a},r}$ and $g_{r,\hat{a}}$ to remove any reference to t_k^r, and removing any unordered pairs in $L_{A\times A}$ that contain $\hat{a}$. The pruning terminates when $L_{A\times A}$ is empty. This step removes any inconsistency in the associations by greedily retaining only those of a certain quality (given by the dissimilarity measures).

V. Experimental Environment

As this paper seeks to address unresolved issues in [7], we make use of the same experimental setup (Fig. 2). We use a virtual wooded environment approximately 700 meters by 500 meters in area, with roads occurring between clusters of trees. The environment has maximum difference in elevation at most 22 meters. Light grey areas reflect locations of sensor platforms.

We make use of four sensor platforms indicated by arrows in Fig. 2. Collectively, these platforms contain three passive RF receivers (indicated by blue arrows), allowing for TDOA-based triangulation, and two EO cameras (indicated by grey arrows). All sensors are located between 5 and 30 meters above the surface. Our experiments involve two vehicles of different color and RF emitting frequencies so as to be distinguishable via standard techniques. The vehicles move exclusively along the gold path, with starting and ending locations varying per experiment.

All experiments make use of co-simulation between Gazebo and Simulink. During the simulation, the vehicles move along some subset of the yellow path in Gazebo, yielding simulated sensor data for each platform. Simulink reads these data, which we process and exploit to push commands from Simulink to manipulate the sensor platforms in Gazebo.

VI. Simulation Results

A. Experimental Setup and Methodology

The scenario of interest consists of two vehicles moving along the yellow path shown in Fig. 2, each starting at opposite ends of the path. The geolocations of each vehicle are known (a feasible assumption since they emit separable RF signals) and available to all platforms. Both EO cameras record video at 10 frames per second at resolution 320 x 240, and have the same fields of view (34.38 degrees horizontal, 26.1 vertical).

Our primary focus in this paper is on the platforms equipped with EO cameras, which we assume are equipped with object detectors that output bounding boxes. Using prior notation and Fig. 2 labeling, at a given time, each target t_j^i that camera i detects is represented as the set $S_{EO}^{i,j} \cup S_{RF}^{i,j}$. The $S_{EO}^{i,j}$ consist of the bounding box centroid (c_j^1, c_j^2) and six statistics $\mu_H, \mu_S, \mu_V, \sigma_H, \sigma_S, \sigma_V$ derived as follows. We first take the sub-image corresponding to each computed bounding box, scale it to a 20 x 20 image, and compute each pixel's representation in HSV (hue, saturation, value) color geometry coordinates. We then consider the distribution of the H, S, and V channels separately and compute the means μ_H, μ_S, μ_V and standard deviations $\sigma_H, \sigma_S, \sigma_V$. These statistics provide a low dimensional encoding of both the location and appearance of each target suitable for transmission in a network with constrained bandwidth. For each platforms, we let $S_{RF}^{i,j}$ consist of the current 3D RF-derived geolocation along with the corresponding velocity estimated from a 1 second history sampled at 10 Hz. Note that standard triangulation techniques permit association between RF and EO features (given current EO camera pose).

Fig. 2. Labeled aerial view of simulation environment.

Algorithmically, we represent t_j^i of all quantities above listed in a prescribed order common to all platforms.

The platforms possess two OKLSPI agents pre-trained on single vehicle scenarios for two separate pan/tilt tasks, one for each modality (EO and passive RF) of sensed information. Both tasks have rewards R_t^M (where $M \in \{EO, RF\}$ denotes modality) of the form $R_0^M = 0, R_t^M = C^M(V_t^M - V_{t-1}^M)$ for $t > 0$, where C^M is a (per-modality) positive constant. The V_t^M are the per-modality value functions, which depend on the particular targets assigned to a given sensor platform. The EO value function is given by

$$V_t^{EO,i} = \left(1 - \frac{\|(\bar{c}_i^1, \bar{c}_i^2) - (160,120)\|^5}{\|(160,120)\|^5}\right), \tag{7}$$

where $(\bar{c}_i^1, \bar{c}_i^2)$ denotes the mean of the bounding box centroids of all targets detected by platform i at time t. Intuitively, V_t^i is a penalty on centering (consistent with those employed in [2, 3, 7]). The RF value function is defined as follows. Let $\boldsymbol{q_c^i}$ denote the unit quaternion representing the orientation of the camera center of platform i. The RF-derived geolocations can be used to derive unit quaternion orientations for each vehicle with respect to platform i; denote by $\boldsymbol{q_v^i}$ the Fréchet mean of such orientations. Then we define $V_t^{RF,i}$ to be the Riemannian distance between unit quaternions, i.e.:

$$V_t^{RF,i} = \cos^{-1}(2(\boldsymbol{q_c^i} \cdot \boldsymbol{q_v^i})^2 - 1). \tag{8}$$

Following the notation of Section IV, we assume that the between platform transforms are the identity, reflecting the scenario where neither platform has knowledge of the other platform's parameters (such as the gimbal angles). The platforms have the same utilities: (1) the dissimilarity utilities P_i are constant multiples of the standard squared error loss, and (2) the outlier utility $O_i\left(x \middle| t_1^i, \ldots, t_{T_i}^i\right)$ is a constant multiple of the minimum squared error loss between x and the t_j^i.

B. Simulated Association Scenarios

To test both the speed and accuracy of our game theoretic Triple-C association methodology, we use geolocations in tandem with simulated EO data. We first randomly orient each cameras field of view to center on a randomly chosen location along the yellow path in Fig. 2. If the camera fields of view overlap sufficiently, we place vehicles randomly along the path in the overlap region We move the vehicles in opposite directions along the yellow path to generate geolocation history for each vehicle. Sufficient vehicle separation in image space allows reliable detection of both vehicles. This makes the scenario admissible for assignment game playing purposes We then generate simulated detections as follows..

EO data simulation is necessary because Gazebo has limited options for depicting (vehicle) models We construct simulated distributions of HSV values by hierarchical sampling. Since we are scaling all detections to a 20 x 20 image, we first segment a square into different regions reflecting the distribution of possible backgrounds: e.g., roads, other land, sky, forest/greenery, etc. After this initial segmentation, we create an additional rectangular foreground segment (essentially a bounding box) representing each vehicle, factoring in the vehicle center and approximate height and width. We label pixels that fall within this rectangular region as corresponding to the vehicle; all other pixels keep their previous background labeling. Because the statistics of interest are shape agnostic, rough segmentations are sufficient. We assign HSV values to each pixel in each region by sampling distributions that reflect color values of each object. For example, we sample HSV values for the vehicle region from distributions reflecting coloration of tires, the vehicle body, and windows. We perform a check to determine if one vehicle would partially occlude another in the image plane. If so, we appropriately blend individual vehicle characteristics into the distributions that generate simulated images. For background features, we sample from distributions consisting of possible, land, vegetation, and sky colors.

To test Triple-C, we use the SAP method described in Section III.A and consider a number of different *annealing profiles* (i.e., choices of sequences of β values in Equation 3). Optimal annealing profiles achieve quick convergence to correct association. We randomly generate 1000 association scenarios as described above. We measure Triple-C performance via four different annealing profiles: one proportional to iteration number, one proportional to the quadratic of iteration number, and the other two twice each of the first two profiles. All games last up to 40 iterations, equal to 4 seconds in our scenario. Figure 3 displays our results. Though both linear profiles outperform the quadratic profiles, likely due to lower β values allowing more exploration, all profiles achieve quick convergence to optimality on average, suggesting that SAP will quickly and accurately compute associations for many types of increasing annealing profiles, suggesting that substantial tuning is not necessary to achieve good associations quickly.

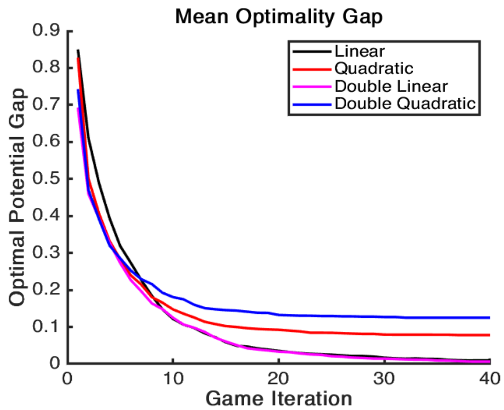

Fig. 3. Per iteration optimality gap for our synthetic association scenarios with SAP under different annealing profiles.

C. Sequential Games in a Target Exchange Scenario

We evaluate the efficacy of our Triple-C method by incorporating it into a closed loop pipeline monitoring a red vehicle and a black vehicle traveling in opposite directions at roughly constant speed along the yellow path in Figure 2, with the red vehicle initially closest to Camera 1 and the black vehicle closest to Camera 2. The OKLSPI policies (described in Section VI.A) control camera pan and tilt. The entire scenario occurs over a 20 second period.

The request and accept game criteria for each sensor platform are the same: request and accept games if multiple concurrent vehicle detections occur in their camera frames. As these criteria are the same for each camera, we consider two separate instances that differ only in which camera requests the game, so as to change the constrain and correct computations in the association process. Once the assignment game starts, we allocate two seconds (20 frames) for the association step, and the subsequent four sections (40 frames) for the actual assignment game. We use track continuity statistics (number of discontinuities and length of discontinuities) as our performance metric, and compare Triple-C performance to two traditional policies for gimbal camera control: fixed position and a uniform velocity sweep along the road. As our past results [2, 3] have shown the superiority of our OKLSPI methodology to other modern deep learning-based camera control methods, we did not include other control policies in our comparison. Given the simplicity of our association scenario, we opted not to test other association algorithms, deferring that for future work where more targets are present in a scenario. Table I presents results from these policies processing the two-vehicle scenario. Neither Triple-C game theoretic assignment policy suffers a track discontinuity, whereas both traditional policies have multiple second periods for which neither vehicle is in view. Moreover, the camera platforms switch custody of the vehicle they are respectively responsible for, with each platform electing to focus on the vehicle closest to it. This behavior is shown in Fig. 4.

TABLE I. TRACK CONTINUITY STATISTICS FOR TRIPLE-C AND TRADITIONAL CAMERA CONTROL POLICIES.

Grid Controller Policy	**Fixed**	**Uniform Sweep**	**Triple-C (Either Camera Start)**
Red # of Discontinuities	3	1	**0**
Blue # of Discontinuities	3	1	**0**
Ave. Red Discontinuity Length (s)	2.50	3.30	**0.00**
Ave. Blue Discontinuity Length (s)	2.27	3.40	**0.00**

As a test of the constrain and correction steps, we repeated the above scenario multiple times, but injected spurious vehicle detections. For example, the requesting node initialized an assignment game by including a vehicle in a region where none is actually present. We limited spurious detections to regions containing a sufficient amount of the ground plane (i.e., no false detections in regions solely consisting of the sky). We assigned no geolocations to spurious detections, making all reported RF statistics zero. In all cases, the spurious detection was either filtered out during the constrain step or not placed in correspondence during the correction step, showing robustness to false detections.

VII. Conclusion

In this paper, we considered a specific multimodal, multi-node association problem that originated out of the assignment game framework presented in [7]. Using the centralized communication structure present in the problem, we proposed the Triple-C method for solving this problem in an intuitive, parallelizable way via game theory. We provided theoretical performance guarantees and showed through multiple experiments that our methodology can quickly and accurately solve the assignment game association problem in realistic scenarios, requiring nodes to communicate minimal information about sensed objects in their purview with no communication of intrinsic sensor parameters required.

There are multiple directions of future work. Scalability of the collective closed loop pipeline of RL + Triple-C + Assignment Games is perhaps of greatest interest, as we have only tested our methods on scenarios with relatively few objects of interest. To that end, we believe that investigating other game theoretic methods outside of the potential game framework we adopted could prove fruitful. Finally, we plan to continue to test our methodology on edge hardware in real environments.

Acknowledgment

The authors thank the Air Force Research Laboratory's Information Directorate for funding this project under AFRL contract FA8750-22-C-1005. The views an conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the United States

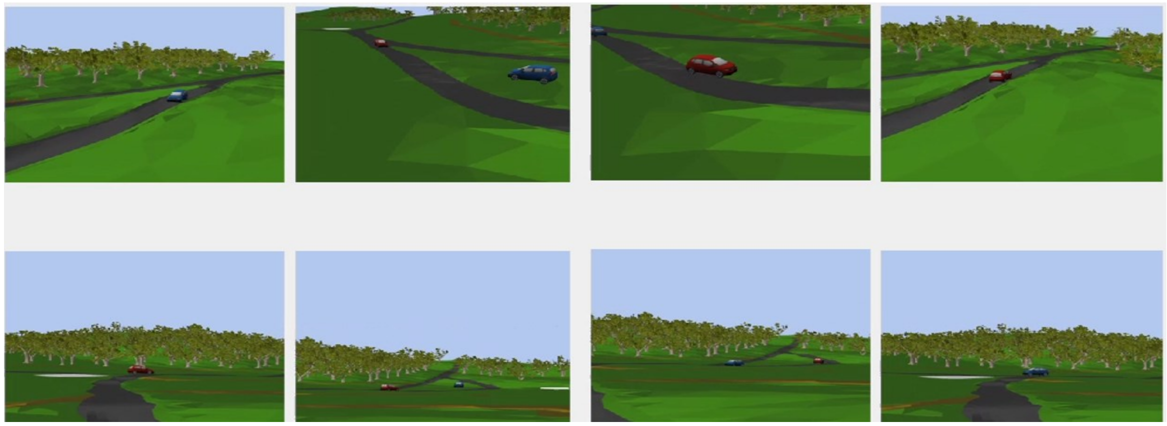

Fig. 4. Triple-C closed loop camera control policy 4, 8, 12, &16 seconds into a 20 second scenario (one camera per row).

Air Force. This paper has been cleared Distribution A, Public Release, under AFRL Case Number AFRL-2024-2953.

REFERENCES

[1] D. Garagić, R. J. Ravier, J. Peskoe, T. Galoppo and P. Zulch, "Submodular optimization via reinforcement learning for active control of sensor networks," in *IEEE Aerospace Conference*, 2022.

[2] R. J. Ravier, D. Garagić, J. Peskoe, T. Galoppo, J. Tigues, B. J. Rhodes and P. Zulch, "Online Reinforcement Learning for Autonomous Sensor Control," in *IEEE Aerospace Conference*, Big Sky, MT, 2023.

[3] R. J. Ravier, D. Garagić, T. Galoppo, J. Peskoe, B. J. Rhodes and P. Zulch, "Reinforcement Learning-based Autonomous Sensor Control via Simultaneous Learning of Policies and State-Action Spaces," in *International Conference on Information Fusion*, Charleston, SC., 2023.

[4] D. Garagić, J. Pesoke, F. Liu, M. S. Claffey, P. Bendich, J. Hineman, N. Borggren, J. Harer, P. Zulch and B. J. Rhodes, "Upstream fusion of multiple sensing modalities using machine learning and topological analysis: an initial exploration," in *IEEE Aerospace Conference*, 2018.

[5] D. Garagić, G. Von Pless, R. Hagan, F. Liu, J. Peskoe, P. Zulch and B. J. Rhodes, "Unsupervised upstream fusion of multiple sensing modalities using dynamic deep directional-unit networks for event behavior characterization.," in *IEEE Aerospace Conference*, 2019.

[6] D. Garagić, D. Pelgrift, J. Peskoe, R. D. Hagan, P. Zulch and B. J. Rhodes, "Machine Learning Multi-Modality Fusion Approaches Outperform Single-Modality & Traditional Approaches," in *2021 IEEE Aerospace Conference*, 2021.

[7] R. Ravier, D. Garagić, T. Galoppo, B. J. Rhodes and P. Zulch, "Multiagent Reinforcement Learning and Game-Theoretic Optimization for Autonomous Sensor Control," in *IEEE Aerospace Conference*, Big Sky, MT, 2024.

[8] R. Geraldes, A. Goncalves, T. Lai, M. Villerabel, W. Deng, A. Salta, K. Nakayama, Y. Matsuo and H. Prendinger, "UAV-based Situational Awareness System Using Deep Learning," *IEEE Access,* vol. 7, pp. 122583-122594, 2019.

[9] K. Anjum, V. Sadhu and D. Pompili, "Multi-UAV Situational Awareness via Distributed and Approximate Computing Techniques," in *IEEE 17th Conference on Mobile Ad Hoc and Sensor Systems (MASS)*, 2020.

[10] T.-Y. Fan, F. Liu, J.-W. Fang, N. Venkatasubramanian and C.-H. Hsu, "Enhancing Situational Awareness with Adaptive Firefighting Drones: Leveraging Diverse Media Types and Classifiers," in *Proceedings of the 13th ACM Multimedia Systems Conference*, 2022.

[11] Y. Bar-Shalom and T. E. Fortmann, Tracking and data association, Acoustical Society of America, 1990.

[12] Y. Bar-Shalom and X.-R. Li, Multitarget-multisensor tracking: principles and techniques, Storrs, CT: YBs, 1995.

[13] E. Blasch and L. Hong, "Data association through fusion of target track and identification sets," in *International Conference on Information Fusion*, 2000.

[14] T. I. Amosa, P. Sebastian, L. I. Izhar, O. Ibrahim, L. S. Ayinla, A. A. Bahashwan, A. Bala and Y. A. Samaila, "Multi-camera multi-object tracking: a review of current trends and future advances," *Neurocomputing,* vol. 552, 2023.

[15] J. Wan and L. Li, "Distributed optimization for global data association in non-overlapping camera networks," in *International COnference on Distributed Smart Cameras*, 2013.

[16] E. Ristani and C. Tomasi, "Features for multi-target multi-camera tracking and re-identification," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2018.

[17] S. Yang, F. Ding, P. Li and S. Hu, "Distributed multi-camera multi-target association for real-time tracking," *Scientific Reports,* vol. 12, no. 1, 2022.

[18] S. Khan and M. Shah, "Consistent labeling of tracked objects in multiple cameras with overlapping fields of view," *IEEE Transactions on Pattern Analysis and Machine Intelligence,* vol. 25, no. 10, pp. 1355-1360, 2003.

[19] N. Jiang, S. Bai, Y. Xu, C. Xing, Z. Zhou and W. Wu, "Online inter-camera trajectory association exploiting person re-identification and camera topology," in *Proceedings of the 26th ACM International Conference on Multimedia*, 2018.

[20] K. G. Quach, P. Nguyen, H. Le, T.-D. Truong, C. N. Duong, M.-T. Tran and K. Luu, "Dyglip: A dynamic graph model with link prediction for accurate multi-camera multiple object tracking," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021.

[21] M. Hofmann, D. Wolf and G. Rigoll, "Hypergraphs for joint multi-view reconstruction and multi-object tracking," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2013.

[22] Y. He, X. Wei, X. Hong, W. Shi and Y. Gong, "Multi-target multi-camera tracking by tracklet-to-target assignment," *IEEE Transactions on Image Processing,* vol. 29, pp. 5191-5205, 2020.

[23] G. Arslan, J. R. Marden and J. S. Shamma, "Autonomous Vehicle-Target Assignment: A Game-Theoretical Formulation," *Journal of Dynamic Systems, Measurement, and Control,* vol. 129, no. 5, pp. 584-596, 2007.

[24] J. R. Marden, G. Arslan and J. S. Shamma, "Cooperative Control and Potential Games," *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics),* vol. 39, no. 6, pp. 1393-1407, 2007.

[25] S. Kirkpatrick, C. D. Gelatt Jr. and M. P. Vecchi, "Optimization by Simulated Annealing," *Science,* vol. 220, no. 4598, pp. 671-680, 1983.

[26] The Univalent Foundations Program, Homotopy Type Theory: Univalent Foundations of Mathematics, Institute for Advanced Study: https://homotopytypetheory.org/book, 2013.