2011 American Control Conference
on O'Farrell Street, San Francisco, CA, USA
June 29 - July 01, 2011

Decentralized Swarming Beliefs of Distributed Autonomous Heterogeneous System

Jamahl Overstreet
Dept. of Mechanical & Aerospace Engineering
Polytechnic Institute of NYU
Six MetroTech Center
Brooklyn, NY 11201

Farshad Khorrami   Prashanth Krishnamurthy
Dept. of Electrical and Computer Engineering
Polytechnic Institute of NYU
Six MetroTech Center
Brooklyn, NY 11201

Abstract— **This paper outlines a novel method for Cooperative Behavioral Control of distributed heterogeneous autonomous systems, emulating the methods in which humans collaborate. This method allows autonomous systems to collaborate on tasks and mission goals, similar to how humans interact, and is effective and efficient for real–time resource allocation. The proposed method fundamentally reduces the required communication bandwidths by significantly decreasing the amount of data necessary for real–time information exchange between cooperating agents. This is done by creating a swarm to estimate the beliefs of the collective, and not on physical states which is usually done by classical approaches. In sharing core beliefs, a collective of heterogeneous agents can plan as an individual, inherently and naturally deconflicting the notion of cost and optimality.**

I. Introduction

In this paper we are looking to explore new techniques to address the fundamental problems of cooperative behavior control. Many contemporary techniques are inadequate since they look to only control the positional states of a collective of agents. With regards to the increasing levels of autonomy for mobile robotic systems, there needs to be a more in depth study of how to utilize the intelligence of these systems for cooperative purposes.

Currently, there is a "rush to market" for many different types of autonomous systems. These technologies, finally reaching maturity, are becoming reliable and very practical to handle sets of objectives with marginal human interactions. As these systems become mass produced, further work needs to be done to find a robust and efficient framework that effectively promotes collaboration amongst systems containing mixed levels of autonomy, for mixed initiative mission sets. There needs to be methods that allow these systems to extrinsically and intrinsically collaborate and cooperate on tasks and mission goals, similar to how humans interact. Further, these methods must be light-weight with respect to CPU processing requirements (since computer resources for most systems are fined tuned to their individual basic requirements), yet effective enough to properly and efficiently coordinate efforts in real–time.

Autonomous systems, artificial or biological, internally utilize large bandwidths of data to coordinate causal actions based upon environmental stimuli. This is due to signals being transported along hard–wired processing components, similar to the way nerves route sensory information to the brain, or wires transmitting data to central processors. When autonomous entities plan amongst each other, they require wireless means of communication, such as sound, radio waves, and stigmergic signals (e.g., hand signals). Wireless negotiating processes create limitations in routing data since all other autonomous entities share the same medium, creating an environment where communication bands collide, limiting and/or disrupting the communication amongst neighboring agents. A realistic solution should mimic nature, which implicitly minimizes the reliance on continual, non–yielding "wireless" communications.

Historically, many works on Cooperative Behavior Control (CBC) deal with abstract robotic systems that contain limited, low fidelity models for analysis and control. Methods for developing CBC architectures and algorithms historically draw upon biological models for inspiration. Most research is guided by ethology, the study of primal animal behaviors, which were used as a basis for many analysis and design models. These primal behaviors were split into primarily three different categories such as swarming [1]–[3] (where individual behaviors are affected and influenced by community behaviors), flocking [4]–[6](a subset of swarming, navigates entities in clusters), and foraging [7], [8] (cooperation to augment individual perception by that of other members in their group).

Currently, cooperative research is expanding with the use of different models and tool sets for analyzing and developing cooperative behaviors. As an example, Martinez is using system theory to analyze emergent behaviors in animal groups and is designing autonomous robotic networks based on this methodology [9]. Caprari is looking at forming controls over mixed societies of robots and animals using control theories [10]. Yanfei and Passino are characterizing swarm cohesiveness as a stability property and use a Lyapunov approach to develop conditions under which local agent actions will lead to cohesive foraging [3]. Dioubate and Mohamed are looking at artificial immune system that imitate the natural immune system [11]. Andrews and Passino are using an evolutionary game-theoretic framework to explain why sociality may emerge in some environments and for some agent objectives [12]. There are also works using a consensus model [13], [14].

Partial funding for this research was provided by CUNY AGEP at Polytechnic Institute of NYU (Award Number 0450360).

Other studies deal with market based approaches, expanding on the concept of economic models [15]. Stentz et al. have developed upon market based methods [16], [17], as well as Viguria et. al. [18], [19]. These approaches typically require large bandwidths, in addition to a common reference for cost.

This study address the cooperation problem by developing a method that learns how other agents in a collective conceive cost, and to merge all notions of cost to a central concept of cost. This is done by each agent sharing their plans amongst the collective, and in turn, each agent in the collective marginally alters their internal perception to align with those of the collective.

II. Motivation

It is known that systems perform much more efficiently when they share the same beliefs. This model is evident in corporations, social networks, institutions, military forces, families, and all other forms of cooperative interactions. An interesting, and a profound point, is that all of the aforementioned group entities behave (locally) as individuals, with individual desires, ambitions, and goals. Each individual is autonomous and intelligent, with vast years of cognitive experience in simulating judgment calls and their respective expected outcome. But, how do all of these individuals, with powerful mental capacities, act and behave similarly in social networks, as an individual organism? These individuals are able to cooperate as an organism, when they share the same beliefs.

We approach CBC with the introduction of an Internal Belief Systems that comprise of the "weights" used to estimate the effort it would take for agents to perform tasks. This technique implicitly addresses the fundamental limitations of communication bandwidths, as well as the deconfliction problem. Note that goals are comprised of a sequence of tasks assigned to agents that satisfy an overall objective.

The method proposed in this study inherently minimizes bandwidth requirements since it models human forms of communication [20]. we are using the AC^{3}E (Autonomous Collaborative Coordinative and Cooperative Environment) construct to develop this methodology. The AC^{3}E framework, as seen in Figure 1, is a higher abstraction then the contemporary notion of swarming. The swarming introduced in this study is broader and more complex than the simplistic spatial alignments of moving entities, but of the swarming of their Internal Belief Systems.

Humans' (who have very complex biological cognitive processes) behavior in masses are also governed by swarming principles. Humans inherently create plans for performing goals. The costs of these plans are based upon each individuals notion of their reality, which is a function of their perception and previous experiences. The perception and knowledge base of experiences comprises of each individuals' belief system.

Similar to humans, each mobile agent will determine an optimal plan for the collective. Their respective "optimal" plan is a function of the "environmental" states (e.g., terrain, weather, relative distances between agents, relative distance between agents and tasks) and their understanding of how costly it would be for a specific agent to perform a specific task, subject to those environmental states. By converging the overall beliefs of each agent, the agents' notions of optimality will be equivalent, and hence they will share a singular notion of global optimality.

With equivalent notions of optimality, there would be no need for an explicit exchanging in cost (e.g., the market based approach). The market based methodology assumes there is a global cost that governs the behaviors of all agents. There is no guarantee that the internal cost for one agent to perform a task is equivalent for another. As an example, say we have an Unmanned Ground Vehicle (UGV) built by two different companies. Company A designs its UGV's utility function to calculate cost in meters, while Company B uses inches. Now both companies' UGVs calculate the cost of a task to traveling between two points as the distance squared. Therefore, Company B's UGV will always exchange the cheapest cost for traveling between two points, which is obviously incorrect. This means, that by using the market based approach, agents would trade cost in different "currency", where as an example, UGV A uses dollars and UGV B uses yen. With respect to potential modifications in cost calculations (due to learning or up–grades), it would be difficult to assign a fixed "exchange rate", further reducing the effectiveness of the market based approach.

Our method concedes that each heterogeneous agent will internally have different methods and weights for determining their own individual costs to perform tasks, which are not shared uniformly across the collective. By implementing a mechanism to allow agents in the collective to converge to a centralized belief system, each agent would effectively share a single "mind". This means that each agent would share a singular notion of optimality, and plan equivalently. In this ideal scenario, no communication would be needed, effectively reducing the required communication to zero. The market approach would not be necessary since each agent would reach the same conclusion. Granted, any variation in the utility functions used to calculate cost would skew each individual notion of optimality from the center. These effects are being analyzed, and will be presented in future works.

III. Proposed Approach

The primary intent of this study is to frame the premise for the Swarm Belief hypothesis. To determine its feasibility, basic assumptions are being considered, (1) each agent knows its own costs, (2) each agent begins with an approximation of its neighbors' cost, and (3) each agent shares the same "currency" (measure of cost). Initially, we are looking to determine if this method is a solution for the relaxed requirements (common currency). The results found in Section VI shows favorable results, demonstrating that this method is a viable method for cooperative behavior control. We are extending these findings to address the convergence of internal beliefs *without* a shared currency. This will also be presented in future work.

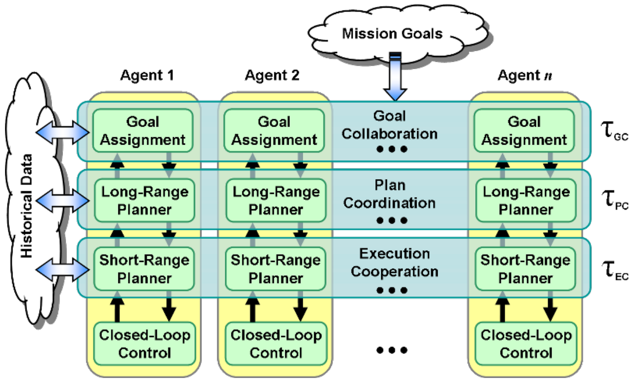

Fig. 1. AC^{3}E Construct

Figure 1 is an example of a basic AC^{3}E construct for collectives of autonomous and semi-autonomous systems to work together. The AC^{3}E model is a derivative of the 4D/RCS model ([21], [22]), where this construct further extends and expands cooperative behaviors. There are primarily four phases in autonomous planning and control for an individual intelligent mobile platform, *(i)* goal assessment, *(ii)* long–range planning, *(iii)* short–range planning, and *(iv)* closed–loop control. Note that goals comprise of a sequence of tasks.

The CBC methods outlined in this work are implemented during the goal assignment phase. To put the implementation into context regarding the overall AC^{3}E paradigm, the complete cooperative cycle is noted. After the goals have been selected, the collective refines their individual intent during the coordination phase of their planning. This is done to make their actions more efficient to accomplish their individually assigned goals, and the group's overall goal. This allows the collective to locally modify the original plan due to unobserved environmental influences that affect the expected state of the collective without having to re-collaborate. During the execution of the plans (i.e., the short–range planning phase), the collective will modify their behaviors to cooperatively accomplish their individual assigned short–range tasks.

The belief system governs the behaviors of each autonomous systems through a series of utility functions that penalizes bad behavior and rewards good behavior. These series of utility functions are usually a function of the magnitude of an environmental influence, and the relative weight assigned to that influence.

For a collective to work efficiently, they need to share a similar belief system. This is difficult when working with heterogeneous systems, since they inherently have different ways of dealing with environmental influences. Systems that learn based upon previous experiences are also difficult to coordinate since they will dynamically change the way they interpret the world, making their beliefs constantly in flux.

Most methods require agents to broadcast their costs for performing tasks in real–time. This method has many deficiencies since *(i)* they need large bandwidths for collaborating, *(ii)* the solutions have a great tendency to fall into a local minima, *(iii)* it is difficult to deconflict costs, *(iv)* most studies use homogeneous agents, where all agents have the same a priori belief of costs and same utility functions, or *(V)* have exotic marketing trade based approaches rooted in economic theory, which will be difficult to realize in real–world constraints.

IV. Generic Belief Structure

The principle behind creating a generic belief structure is the understanding that all utility functions share a common property: (*i*) they estimate cost by determining the expected effort an agent would exert to negotiate an environment, based upon the type and state of that environment, (*ii*) the state of the tasks, and (*iii*) the state of the team members. A generic utility function can be described as a product of weighted sums.

The generic belief structure has the following form:

$$C^A\left(T_k\right) = \sum_H \sum_E \sum_{i \in T_k} W_H^A \cdot {}^H W_E^A \cdot {}^E X_i^A, \tag{1}$$

$$\bar{C}^A = \sum_k C^A\left(T_k\right), \quad \bar{C} = \sum_{A=1}^{N_A} \bar{C}^A, \tag{2}$$

where, T_k is a specific task; $C^A\left(T_k\right)$ is the cost of agent A performing task T_k; $\bar{C}^A$ is the total cost of agent A performing all tasks; $\bar{C}$ is the sum of component cost for each agent performing assigned tasks; N_A is the number of agents in the plan; W_H^A is the weight of a specific category of environmental influences (H) relative to other categories for agent (A); ${}^H W_E^A$ is the weight of a specific attribute (E) of the specific environmental category (H) for agent (A); ${}^E X_i^A$ is the state (i) of the attribute (E) for agent (A).

The capital letters (e.g., ${}^E X_i^A$) signify the generic classification designation (e.g., such as W_H^A specifying a terrain type vs. specifying a weather condition) of the weights and states. Lower case letters signify the specific classification (i.e., w_m is the cost of traversing over mud). The two weights (i.e., W_H^A and ${}^H W_E^A$) are used to balance the relative effect each weight has on the classification of environmental influences. This is necessary to combine costs of mutually exclusive quantities relative to their respective categories. As an example, mixing the cost of performing a task may be based on a different unit of measure than that of the cost of navigating over a specific terrain. Therefore, the elemental weight factor (${}^H W_E^A$) determines the relative cost of each element within a category, and the category weights (W_H^A) determines the relative costs between different categories.

A. Belief Convergence

Initially, we consider a basic generic method for emerging intelligence, where agents will look at each individual plan and adjust their belief based upon their neighbors' best plans. Each individual agent will preform an error back–propagation gradient–descent of the relative errors amongst the neighboring agents' estimated cost relative to their own individual cost for a specified plan. After a finite number of training cycles, all of the agents' beliefs should converge close to the weights and costs of the collective.

The weight update has the following form:

$$\left(W_k^A\right)' = W_k^A + \Delta W_k^A, \tag{3}$$

where W_k^A is a vector, and

$$W_K^A = \bigcup \left\{ W_X^A \in \mathbb{R}^M, {}^X W_E^A \in \mathbb{R}^{m_i} \right\}. \tag{4}$$

For each individual weight,

$$\Delta W_k^A \triangleq -\eta \nabla \left. \varepsilon^A \right|_{W_k^A}, \text{where,} \quad \varepsilon^A = \frac{1}{2}\left(\mathbb{E}(c) - \tilde{c}\right)^2, \tag{5}$$

and where η is the rate of learning, ε^A is the squared error between the estimated ($\tilde{c}$) and expected ($\mathbb{E}(c)$) costs. $\mathbb{E}(c) = \sum_{a=1}^{N_a} \mathbf{P}\left(c^a = c^*\right) c^a$, and N_a is the number of agents. In Equation 5, the derivatives are being taken with respect to W_K^A. The influence of a specific weight on the error is formulated as (Note the lower case w is used to specify a specific weight):

$$\nabla \left. \varepsilon^a \right|_{w_k^a} = \frac{\partial \varepsilon^a}{\partial w_k^a} = \frac{\partial \varepsilon^a}{\partial \tilde{c}} \frac{\partial \tilde{c}}{\partial w_k^a}, \quad \text{and} \tag{6}$$

$$\frac{\partial \varepsilon^a}{\partial \tilde{c}} = \left(\mathbb{E}(c) - \tilde{c}\right), \tag{7}$$

$$\frac{\partial \tilde{c}}{\partial w_k^a} = \sum_{i \neq k \in W^K} \prod_{j \in \bar{K}} w_i^j x_i, \tag{8}$$

where $\bar{K}$ is the set of weights that are not contained in the class of weights to which the derivative is taken. The updated weight has the following generic form:

$$\left(w_k^a\right)' = w_k^a - \eta \left(\mathbb{E}(c) - \tilde{c}\right) \sum_{i \neq k \in W^K} \prod_{j \in \bar{K}} w_i^j x_i. \tag{9}$$

B. Example Belief Swarming: Setup and Training

As an example, suppose a Unmanned Ground Vehicle (UGV) agent is given an assignment of a neighboring agent performing an Information Surveillance, Reconnaissance (ISR) task. The task is influenced by the following environmental considerations, (*i*) the agent transitioning to the task, (*ii*) the condition of the terrain it must maneuver over, (*iii*) weather conditions, (*iv*) the cost of the agent performing the given task, and (*v*) the cost of the agent itself. The following is an example utility function that describes an agent's internal belief[1], using Equation 1, is

$$c_i^a = w_\beta^a \Bigg(\underbrace{w_{Tr}^a \left(w_d^a x_i^d + w_{ms}^a x_i^{ms}\right)}_{\text{Transition Cost}} + \underbrace{w_G^a \left(w_m^a x_i^m + w_g^a x_i^g\right)}_{\text{Ground Cost}} + \underbrace{w_W^a w_s^a x_i^s}_{\text{Weather Cost}} + \underbrace{w_T^a w_{isr} x_i^{isr}}_{\text{Task Cost}} + \underbrace{w_v^a}_{\text{Platform Cost}} \Bigg), \tag{10}$$

[1]Note that the leading category notation for the attribute weights are removed for brevity.

where c_i^a is the overall cost for agent "a" to perform task "i" (the ISR task). w_β^a is the confidence that the agent has of agent a's belief in the world. w_{Tr}^a is the scaling weight that adjusts how costly transitioning to the task is relative to other categories of environmental influences. This is important as to proportionalize the effect of one category of costs over another since some costs may be mutually exclusive and a judgement has to be made to determine how much more expensive should one cost be over another. w_d^a is the weight used to determine the cost of agent $a$to move from its current location to the location of the task. x_i^d is the distance between the state of agent a and the task. w_{ms}^a is the weight of transition from the agent's current location to the task's location, at a given speed. w_G^a is the weight of the influence that ground mobility has relative to other categories of environmental influences. w_m^a is the weight of traveling over a muddy terrain. x_i^m is the state of the mud (the "muddiness") the vehicle will negotiate. w_g^a is the weight of traveling over grass. x_i^g is the state of the grass. The rest of the weights and states follow the same logical sequence, with the exception to w_v^a, which is the bias weight that expresses the necessity of utilizing a particular agent over another.

V. Simulation Setup

Different case scenarios are used to demonstrate belief swarming and the overall performance of collaborating on missions. The Robot Human Access (RHA) user interface and Collaborative and Cooperative Command and Control (C4) software, developed to study cooperative systems, are used to simulate the collaboration amongst neighboring agents. Each mission will comprise of assigning agents to a series of tasks, either in groups or individually. It is assumed that each agent is completely autonomous and needs no human intervention. Figure 2 shows a sample perceptive world map. During the Collaborative Goal phase of planning, a course discretized grid is used for planning. Lower levels of planning, such as the Long–Range Coordination and Short–Range Execution phases will provide better approximations to feedback to the Collaborative Goal planner. Section III and Figure 1 discusses the planning cycles in more detail.

As an example, in the Long–Range Planning phase the agents will adjust the weights that govern the mobility belief. Therefore, the Collaborative Goal will not consider the optimal mobility cost to reach the task, but will only look at the "best fit" linear distance between the initial location of the agent, and the location of the task, and the environmental states along those points. Figure 2 shows a coarse discretized map that is used for Collaborative Goal planning. Each cell in the grid contains environmental information, as shown in Table I, which outlines the environmental attributes of the terrain.

The terrain cell attributes are quantified by the percentage of environmental influences in each cell. As an example, if cells contain grass, it takes into consideration the amount of actual grass to the total amount of grass that would be present if the cell was completely covered by grass. We will be using

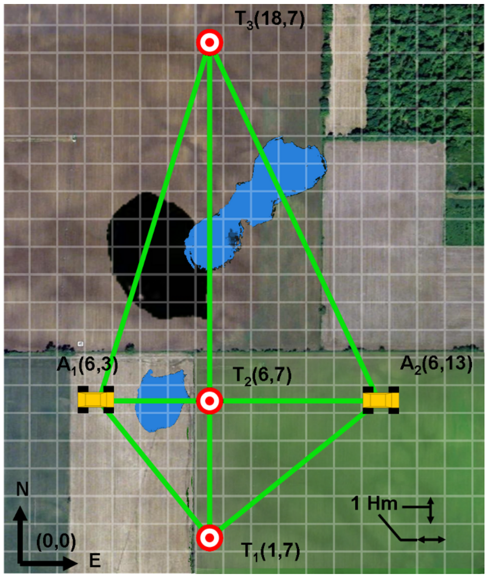

Fig. 2. Sample Perceptive World Model (PWM) Map

TABLE I
SAMPLE OF ATTRIBUTES CORRESPONDING TO FIGURE 2

Cell No	131	132	133	185	186	187
Grass (%)	0	0	39	0	0	0
Mud (%)	0	0	0	0	41	34
Uneven Terrain (%)	8	2	0	0	0	0
Unknown (%)	0	0	0	35	60	0
No Goes (%)	60	23	0	0	0	0
Wind	42	45	37	33	31	34
Wind-Dir	-39	-39	-39	-30	-30	-30
Blowing Storm	1	1	1	1	1	1
Rain	0	0	0	0	0	0

this as a basis for determining how difficult it would be to navigate over terrains of specific attributes. As an example, if an agent has difficulties driving in mud most likely the agent can find areas to maneuver around the mud if the entire region has very little mud. If the region is 80% covered in mud, then the cost will be reflective of this, since the vehicle is more likely to get stuck in the mud.

A. Case Scenarios

In this paper, three scenarios were evaluated and only the the attribute weight of the generic belief utility function was used. Each scenario was developed and analyzed looking at the overall collaborative behavior of agents forming a consensus. The scenarios represent (*i*) 2 agents assigned to 3 tasks, (*ii*) 3 agents assigned to 4 tasks, and (*iii*) 4 agents assigned 5 tasks. This was done to simulate an asymmetric assignment situation, where at least one agent will be assigned to two tasks. The assumptions made are (*i*) each agent's initial belief estimates of their neighbors' belief states are within $1{-}\sigma$ (i.e., 10

A sample scenario is represented by Figure 2, where there are two autonomous heterogeneous agents and three target location. The PWM has mixed terrain types, and each agent type has different strengths and weaknesses, meaning different standard weights for different terrain and weather types. This study looks at only Unmanned Ground Vehicles (UGVs), but can be extended to other types of unmanned vehicles. The collective goal is for all target locations to be visited by any agent.

In the scenario depicted by Figure 2, if all of the agents are homogeneous with equivalent Belief Systems and the terrain is featureless (absence of terrain information, unknowns, no-goes, etc.), the optimal solution will be

$$S^* = \{A_1 : \{T_3\}, A_2 : \{T_1, T_2\}\}. \quad (11)$$

This equation states that the optimal solution is Agent$_1$ performing Task$_3$, and Agent$_2$ performing Task$_1$ and Task$_2$, sequentially.

VI. Results

The results of the different scenarios evaluated are as expected. The environmental influences affect the notion of belief for the different agents. All of the agents' beliefs converged to a central belief. As can be seen in in Figure 3(a) through Figure 3(c), as the the individual beliefs converge close to a singular belief, the agents begin to plan equivalently. The training cycle ends when all agents error norms reach close to steady state. The "Belief Error" shown in the figures are the sum of absolute differences between estimated and actual values for each agent.

In the scenario depicted by Figure 2, the actual optimal solution is $S^* = \{A_1 : \{T_2, T_3\}, A_2 : \{T_1\}\}$. Agent$_1$'s weights were selected to be less sensitive to terrain conditions than than Agent$_2$ weights. Therefore, as in contrast to the optimal solution found in Equation 11, Agent$_1$ is able to handle the terrain (i.e., water, unknowns, uneven) better over time; therefore, it takes over accomplishing task$_2$ and Task$_3$. Agent$_2$ visits only Task$_1$. This scenario demonstrates that the agents' planners are properly interpreting the world. Both agents form the same conclusion of optimality and plan as an individual.

Also as expected, the more the degrees of freedom the search graph contains (the number of nodes opened per child, which is the combination of agents working on tasks), the longer it takes for each Internal Belief Systems to converge to a Central Belief System. This inherently deconflicts the overlapping of assignments, since all agents will have the same belief of optimality and therefore will generate the same plans. If the environment does not change, the agents could potentially continue to perform other missions without any communication.

VII. Conclusion

In this study, we presented a method for giving a group of autonomous heterogeneous systems the capability to deconflict costs and minimize communication bandwidths by converging their internal beliefs. By performing a gradient descent on the respective beliefs of each agent, the team members were able to learn how their neighbors viewed

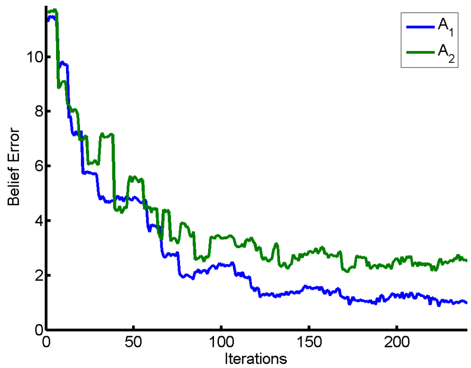

(a) Belief convergence of 2 agents assigned 3 tasks

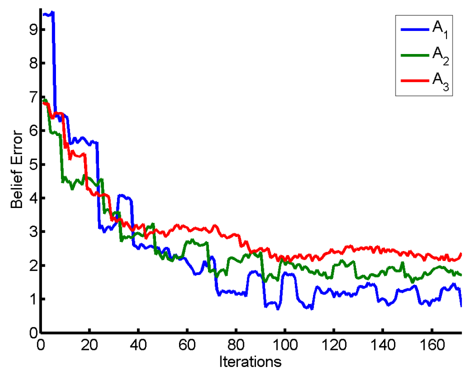

(b) Belief convergence of 3 agents assigned to 4 tasks

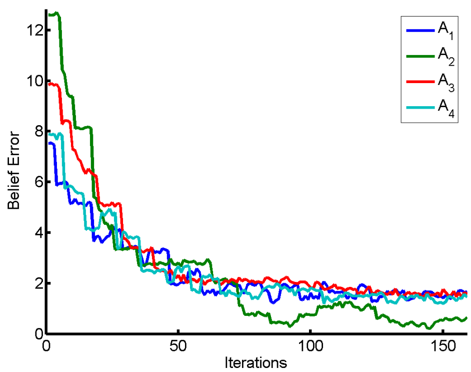

(c) Belief convergence of 4 agents assigned 5 tasks

Fig. 3. Belief Swarming results.

the "world", and then swarm their respective beliefs in alignment. In sharing congruent beliefs, dissimilar agent types can plan as an individual, inherently deconflicting the notion of cost and optimality, while minimizing the necessity for further communications.

REFERENCES

[1] G. Beni and J. Wang, "Theoretical problems for the realization of distributed robotic systems," *Robotics and Automation, 1991. Proceedings., 1991 IEEE International Conference on*, pp. 1914–1920 vol.3, 9-11 Apr 1991.

[2] G. Beni and P. Liang, "Pattern reconfiguration in swarms-convergence of a distributed asynchronous and bounded iterative algorithm," *IEEE Transactions on Robotics and Automation*, vol. 12, pp. 485–490, June 1996.

[3] Y. Liu and K. Passino, "Stable social foraging swarms in a noisy environment," *IEEE Transactions on Automatic Control*, vol. 49, pp. 30 – 44, Jan. 2004.

[4] C. Reynolds, "Flocks, herds and schools: A distributed behavioral model," *Computer Graphics*, vol. 21, no. 4, pp. 25–34, 1987.

[5] N. Moshtagh, A. Jadbabaie, and K. Daniilidis, "Distributed geodesic control laws for flocking of nonholonomic agents," in *Proceedings of the 2005 European Control Conference 44th IEEE Conference on Decision and Control*, pp. 2835–2840, Dec. 2005.

[6] N. Moshtagh, A. Jadbabaie, and K. Daniilidis, "Vision-based control laws for distributed flocking of nonholonomic agents," in *Proceedings of the 2006 IEEE International Conference on Robotics and Automation*, pp. 2769–2774, May 2006.

[7] D. A. Shell and M. J. Mataric, "On foraging strategies for large-scale multi–robot systems," in *Proceedings of the 2006 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 2717–2723, Oct. 2006.

[8] T. Keviczky, F. Borrelli, and G. Balas, "Stability analysis of decentralized rhc for decoupled systems," in *Proceedings of the 2005 European Control Conference 44th IEEE Conference on Decision and Control*, pp. 1689–1694, Dec. 2005.

[9] S. Martinez, J. Cortes, and F. Bullo, "Motion coordination with distributed information," *IEEE Transactions on Control Systems Magazine*, vol. 27, pp. 75 –88, Aug. 2007.

[10] G. Caprari, A. Colot, R. Siegwart, J. Halloy, and J.-L. Deneubourg, "Animal and robot mixed societies: building cooperation between microrobots and cockroaches," *IEEE Transactions on Robotics Automation Magazine*, vol. 12, pp. 58–65, June 2005.

[11] M. Dioubate, T. Guanzheng, and L. Toure Mohamed, "An artificial immune system based multi-agent model and its application to robot cooperation problem," in *Proceedings of the 2008 7th World Congress on Intelligent Control and Automation*, pp. 3033 –3039, June 2008.

[12] B. W. Andrews, K. M. Passino, and T. A. Waite, "Social foraging theory for robust multiagent system design," *IEEE Transactions on Automation Science and Engineering*, vol. 4, pp. 79 –86, Jan. 2007.

[13] D. Castanon and D. Teneketzis, "Further results on the asymptotic agreement problem," *IEEE Transactions on Automatic Control*, vol. 33, pp. 515 –523, June 1988.

[14] Y. Cao and W. Ren, "Optimal linear-consensus algorithms: An lqr perspective," *IEEE Transactions on Systems, Man, and Cybernetics, Part B: Cybernetics*, vol. 40, pp. 819 –830, June 2010.

[15] Y.-L. Chen, B. Gregory, M. Easley, M. Peot, J. Lee, and T. Altshuler, "Market-based collaborations for autonomous operations of unmanned air vehicles," in *Proceedings of the 2006 IEEE Workshop on Distributed Intelligent Systems: Collective Intelligence and Its Applications*, pp. 273 –278, June 2006.

[16] N. Kalra, D. Ferguson, and A. Stentz, "Hoplites: A market-based framework for planned tight coordination in multirobot teams," in *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*, pp. 1170 – 1177, Apr. 2005.

[17] M. Dias, R. Zlot, N. Kalra, and A. Stentz, "Market-based multirobot coordination: A survey and analysis," *Proceedings of the IEEE*, vol. 94, pp. 1257 –1270, July 2006.

[18] A. Viguria and A. Howard, "An integrated approach for achieving multirobot task formations," *IEEE/ASME Transactions on Mechatronics*, vol. 14, pp. 176 –186, Apr. 2009.

[19] A. Viguria, I. Maza, and A. Ollero, "Set: An algorithm for distributed multirobot task allocation with dynamic negotiation based on task subsets," in *Proceedings of the IEEE International Conference on Robotics and Automation*, pp. 3339 –3344, Apr. 2007.

[20] C. L. Butler and A. Rothstein, "On conflict and consensus: A handbook on formal consensus decisionmaking," 2006. http://www.consensus.net.

[21] C. Schlenoff, R. Madhavan, J. Albus, E. Messina, T. Barbera, and S. Balakirsky, "Fusing disparate information within the 4D/RCS architecture," in *Proceedings of the 2005 8th International Conference on Information Fusion*, vol. 2, p. 8, July 2005.

[22] J. Albus, "4–D/RCS reference model architecture for unmanned ground vehicles," in *Proceedings of the 2000 IEEE International Conference on Robotics and Automation*, vol. 4, pp. 3260 –3265, Apr. 2000.