Multi-Robot Belief Propagation for Distributed Robot Allocation

Jonas Nathan Schwertfeger   Odest Chadwicke Jenkins
Department of Computer Science
Brown University
Providence, RI 02912-1910
{js,cjenkins}@cs.brown.edu

***Abstract*— Establishing collaborative behavior is an important factor in coordinating teams of robots. Multi-robot task allocation is one aspect of group coordination that deals with the assignment of robots to subtasks. Toward this end, we propose Multi-Robot Belief Propagation (MRBP), a synthesis of distributed probabilistic inference and notions of theory-of-mind for the purpose of multi-robot task allocation. MRBP does not rely upon a central planner or fixed decision hierarchy, but allows to infer task assignments locally through Bayesian belief propagation. The belief propagation algorithm provides us with a tool to incorporate task-related intentions and beliefs of robots into the assignment process. An assignment process in which each individual robot probabilistically decides, based on its own observations and the intentions and beliefs of other robots, which task it should allocate itself to. As a result, MRBP provides a means for collaborative group behavior without explicit protocols and command hierarchies. We present a sample implementation of MRBP for the chain-of-sight task allocation problem and show the results obtained from physically simulated robot teams.**

***Index Terms*— Multi-robot Task Allocation, Belief Propagation, Theory of Mind**

I. Introduction

Robotic technology is being deployed in greater numbers through mass-production efforts, for entertainment and service, and for realizing the grand ambitions of society, such as space exploration. As the deployment increases, so does the importance of the question on how to coordinate groups of robots to perform common tasks. If such groups are to collaborate autonomously on the group level as well as the individual robot level, then synchronization and communication of some sort must be in place between individual robots. The problem of *multi-robot task allocation* (MRTA) pertains directly to this issue and asks how a collection of N robots should be allocated to a set of M subtasks such that a common goal is achieved. The straightforward approach to MRTA is to use a single *central planner* that compiles all relevant information, generates a plan, and unilaterally allocates robots to subtasks. However, this approach possesses two major weaknesses: First, as the number of robots grows, communication and planning costs grow beyond tractability. Second, a fault in the central planner or the failure of individual robots requires protocols for recovery which can be difficult to implement in practice and are often error-prone.

In light of these issues, it seems natural that an individual robot should be its own decision maker and decide by itself on what subtask it should be allocated. Although human society forms into decision making hierarchies, individual people choose their own pursuits based on their sensibility and the beliefs they have about others. In analogy to this, MRTA should allow for robots to allocate themselves based on their innate behavior and their beliefs about other robots. That way MRTA does not have to rely on the dictates of central planners and can recover from unexpected events in a flexible manner. Additionally, a belief-based method offers the flexibility of allocating with uncertainty, when situational ambiguity prohibits making hard decisions.

Toward these ends, we propose *Multi-Robot Belief Propagation* (MRBP), a probabilistic method for distributed multi-robot task allocation that incorporates notions from the *theory-of-mind* concept [9]. In MRBP we cast the MRTA problem as a distributed statistical inference problem in which each robot computes a probability distribution over the possible subtasks it could be allocated to. These distributions correspond to marginal probabilities of "hidden" variables on a *Markov random field* (MRF) [6] and are approximated in a distributed fashion using the Bayesian *belief propagation* algorithm (BP) [14]. In the MRF model that we present each robot infers this marginal from its own observations of the world and messages based on beliefs of other robots.

We demonstrate MRBP through application to the *chain-of-sight* (COS) problem with physically simulated robot teams. The chain-of-sight problem involves a set of robots arbitrarily distributed in an environment which start to form a chain of visibility between a particular start and goal location. We use chain-of-sight as a sample application to demonstrate the potential for other areas, such as search & rescue operations and the maintenance of communication in mobile mesh networks.

II. Background

Our approach is a synthesis of different perspectives on group coordination that appear in three distinct areas of AI: multi-robot task allocation, distributed probabilistic inference, and cognitive models of development and shared belief. In the following paragraphs we present a brief review of each area and highlight their connections.

Multi-robot task allocation approaches have been proposed in great variety, such as by [11], [1], [8]; however, here we guide our focus to two specific formulations: The optimal assignment problem (OAP) and auctions. [7] cast MRTA as a centralized optimal assignment problem where, given an $M \times N$ utility matrix, combinatorial optimization finds the allocation with the maximal utility. [5] phrase MRTA as a market economy where leaders develop plans for a group of robots and uses them to bid for tasks against other bidding robots at an auction. Since every robot acts as to maximize its utility (or profit), an MRTA solution emerges from the group economy. While these methods address MRTA for heterogeneous robots, they impose a hard phrasing of robot-to-task utility and a rigid leadership structure. Instead, we aim for a soft, probabilistic representation of utility and group inference that utilizes interaction without imposing a decision hierarchy.

Distributed probabilistic inference is related to the problem of constructing a joint probability distribution over a set of latent variables through the pairwise interaction of such variables. Each latent variable maintains a "belief" about its assignment that is conditionally dependent on its own observed variables and the beliefs of its neighboring latent variables. Such a model is a good fit for MRTA in that each robot will maintain a belief about its assignment to a known set of subtasks using its own observations and communications with other robots. To perform inference, we model the variables and their dependencies as a Markov random field and run loopy belief propagation on it. Convergence and general properties of BP have been described by [14]. Yedidia et al. also discuss the tight relationship to Ising models which are used by [13] for describing and coordinating macroscopic group behavior. [10] describe a probabilistic finite state automata approach to MRTA where group behavior emerges implicitly from environmental interaction but no explicit message passing mechanism or group coordination is carried out.

Cognitive models of development attempt to bring understanding of human behavior both as individuals and groups. In terms of developmental robotics, such models can inform the development of robot controllers to function individually and in groups. Specifically, development helps robotics to understand what functions should be innate (preprogrammed), adaptive (learned), and informed by social interactions. We pay specific attention to Leslie's Theory of Mind (ToMM) model [9] that provides a framework for describing how intentionality drives decision making. Such intentionality is embodied by an individual's innate self-direction towards goals and influenced by the perception of shared beliefs with others (as emotions). Although the validity of ToMM in reality cannot be formally stated, it does have a form suitable to integration with MRTA and MRFs. This inspiration was founded on previous works of Scassellati [12] and Breazeal [3] in their application of ToMM to human-robot interaction with humanoid robots. Our work aims to generalize these concepts through incorporation into the probabilistic MRF framework. The expression of innate behavior and shared belief of ToMM has a very convenient breakdown into the message passing formulation on MRFs, which we explain in later sections. Alternative to ToMM, Bratman's Belief-Desire-Intention (BDI) model [2] is another approach to modeling decision making in groups with a greater emphasis on deliberative reasoning and planning.

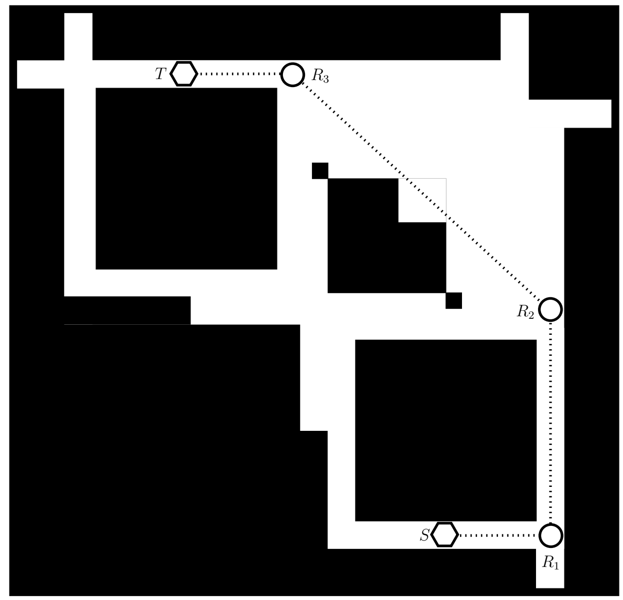

Fig. 1. Chain-of-sight arrangement of three robots on a building floor. The two hexagons represent the goal object and starting location whereas the circles represent the robots.

III. Chain-of-Sight Example

To ground our discussion of MRBP, we consider the chain-of-sight (COS) problem as an example of how MRBP can be utilized. In COS, we assume that multiple robots spread out in an environment where there is a known start location and a goal object at an unknown location. The group of robots searches for the goal object and once found, forms a visibility chain between the start and the goal. Individual robots can communicate within some range and perceive other robots and the goal location through their field of view.

The task allocation problem consists of assigning robots to discrete locations in order to build the visibility chain. We assume there are a fixed enumeration of valid robot locations on the map available for allocation. Fig. 1 shows an instance of this problem with the robots already positioned in a valid configuration. Real-world applications where this kind of behavior could be of interest are ad-hoc formation of communication networks where the sight criterion is substituted by a signal range criterion, or search & rescue missions in a post-disaster scenario. Our implementation assumes that each robot has a map of the terrain and perfect localization on it, but in principal, task allocation could be complementary to simultaneous localization and mapping (SLAM).

IV. Multi-Robot Belief Propagation

A. MRTA as a Statistical Inference Problem

Casting MRTA into MRBP, we consider each robot i to maintain two random variables: y_i, an observed variable representing a robot's own perception, and x_i, a hidden variable representing the probability of allocation to all possible subtasks. Although these variables can either be discrete or continuous, we assume x_i is a discrete enumeration over the set of possible subtasks. For COS, x_i enumerates over locations on the map. The probability of robot i being allocated to certain tasks, which we refer to as the belief $b_i(x_i)$ of i, depends on i's observations y_i and "advice messages" from other robots in its communication neighborhood, where the set of neighbors of robot i is denoted by $N(i)$. Fig. 2 illustrates these variables and dependencies in a graphical model. The belief propagation algorithm allows us to approximate the marginal x_i through passing of advice messages (referred to as *belief messages* by [14]). The belief $b_i(x_i)$ of robot i is then given by the following product:

$$b_i(x_i) = Z\phi_i(x_i, y_i) \prod_{j \in N(i)} m_{j,i}(x_i) \tag{1}$$

where Z is a normalization constant that ensures the belief sums to one. $\phi(x_i, y_i)$ is called the local evidence function and expresses the likelihood of robot i being allocated to a specific subtask based on its own observations y_i. $m_{j,i}(x_i)$ is a belief message from robot j to robot i suggesting how i should be allocated. Each robot j forms such a message to all robots i in their neighborhood $N(j)$ as a product of its local evidence, the compatibility between task allocations, and incoming messages from its neighborhood (except for those coming from i), summed over all possible allocations of x_j:

$$m_{j,i}(x_i) = \sum_{x_j} \phi_j(x_j, y_j)\psi_{j,i}(x_j, x_i) \prod_{k \in N(j) \setminus i} m_{k,j}(x_j) \tag{2}$$

$\psi_{j,i}(x_j, x_i)$ is a compatibility function that expresses how much robot j would recommend a specific allocation of x_i given its own possible allocation x_j.

It has been shown that belief propagation finds exact marginals in graphs without loops and approximations in graphs with loops [15]. Since our model has loops we have to settle for approximations.

B. MRBP and Theory of Mind

To this point, we have only provided a framework for distributed inference on allocation, leaving several terms that still need to be defined. In particular, we need to define the local evidence function $\phi_i(x_i, y_i)$ and compatibility function $\psi_{i,j}(x_i, x_j)$. To define these functions, we take inspirations from ToMM [9] about how innate goal-directed behavior and social communication can inform robot control, such as in uses by Scassellati [12]. An illustration of the cognitive model behind MRBP is presented in Fig. 3.

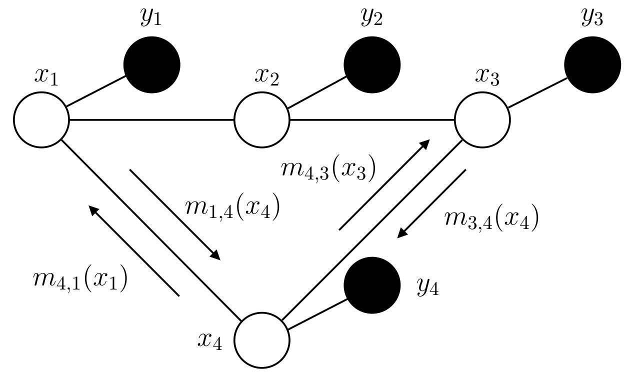

Fig. 2. An instance of the graphical model of MRBP. Robot i's hidden variable x_i depends on the observed variable y_i and the hidden variable x_j of each robot j in i's neighborhood. A message passed from i to j, that is i's belief about j, is depicted by $m_{i,j}(x_j)$. In this graph each robot has two neighbors. In general, however, the actual neighborhood graph is problem-specific.

We cast $\phi_i(x_i, y_i)$ as a robot's innate goal-directed behavior. This function combines all factors that relate the robot's perceptions to its innate individual behavior. In case of COS, this behavior can be formulated as the product of three functions:

$$\phi_i(x_i, y_i) = D(x_i, y_i)G(x_i, y_i)O(x_i, y_i) \tag{3}$$

Although these factors are described at more depth in the COS implementation section, we briefly outline their functionality. The factor $D(x_i, y_i)$ represents traveling distance and expresses a robots preference for closer locations, $G(x_i, y_i)$ stands for goal attraction and assures that a robot does not move if it sees the goal and $O(x_i, y_i)$ expresses occupancy of locations by other robots or obstacles.

We cast $\psi_{i,j}(x_i, x_j)$ as a robot's mechanism for social interaction regarding task allocation. This function allows robot i to evaluate how suitable allocations of x_j are to its own allocations. In the chain-of-sight example, this function expresses a high compatibility between locations that are visible to each other and lie on the shortest path toward the start location.

In sending messages to robot j and making its decisions, robot i could use a notion of shared belief about robot j. We use the term $b_i(x_j)$ to express the belief robot i has about the belief of robot j. Robot i can use this belief to anticipate the actions of robot j, regulate the "insistence" (entropy) of its messages to j, or adjust its own plans. Note, $b_i(x_j)$ is not the actual belief of robot j about itself, but rather robot i's estimation of what j's belief about itself is. In the COS example, we assume shared beliefs are observed exactly, through exchange of the true belief $b_j(x_j)$ between robots, but this need not always be the case. Shared beliefs could be considered latent variables that are learned from prior experience with certain other robots, or inferred over the course of interaction.

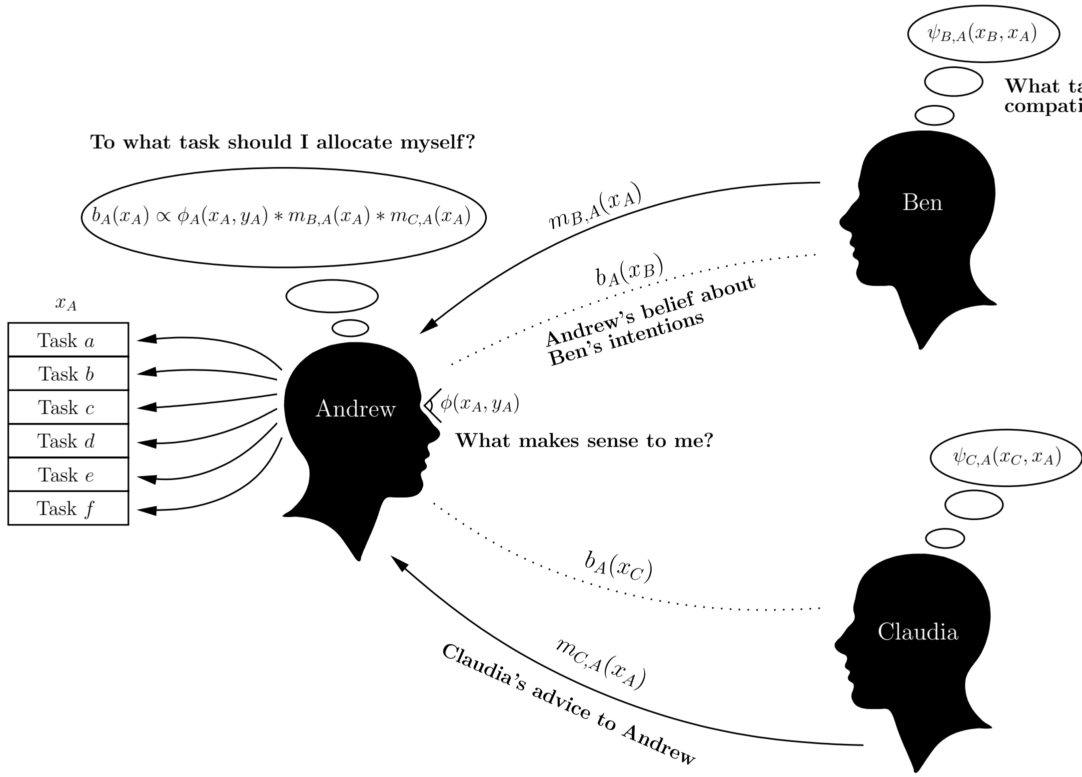

Fig. 3. An illustration of the cognitive model behind MRBP in terms of beliefs, intentions and advice between three humans, centered at "Andrew". Andrew has to decide which task he should allocate himself and forms a belief $b_A(x_A)$ about his allocation. This belief is formed as a combination of what allocation "makes sense" to him ($\phi(x_A, y_A)$), based on its own perception of the world (y_A), and the advice he receives from Ben and Claudia. Ben and Claudia form their advice for Andrew in terms of how beneficial it would be for them if Andrew would allocate himself to certain tasks.

C. Practical Considerations

Thus far, we described the basic MRBP method but have not addressed how robots communicate to each other and how they form neighborhoods. For message passing any form of directed communication is acceptable. In order to be standard compliant and allow integration of heterogeneous robots we suggest the usage of a well-established network layer protocol such as IP. On top of IP the transport protocol UDP is reasonable for message passing if a single message fits into a single datagram. If a message is fragmented into several datagrams and a single datagram is not successfully transmitted, then the message is incomplete and must be disregarded. In these cases TCP might be a better choice. On the lower data link and physical layers the choice of an ad-hoc technology seems more robust to failures and unforeseen events than a fixed infrastructure network. Popular ad-hoc options include wireless local area networks in ad-hoc mode and Bluetooth Piconets.

If an application scenario allows to predefine which robots can and should exchange belief messages between each other, that is if all neighborhoods are known a priori, and unexpected changes can be ruled out, then there is no need to dynamically build neighborhoods. For the COS problem and many other interesting scenarios, however, the relations between robots are not known a priori and vary during execution of tasks; robots move in and out of communication range or they fail to work. Hence, robots have to explore and update their neighborhoods continuously. A natural way to accomplish this consists of periodically broadcasting "beacon" data packets containing the sender's identification and maintaining a list of neighboring robots based on received beacons. If no beacon is received from a robot over a certain period of time it is removed from the neighborhood. By removing a robot from the neighborhood we also discard the last message we received from it and consequently, that robot's belief is not incorporated into Eq. (1) and (2) anymore. If a beacon is received from a robot that is not in the neighborhood, then that robot is added.

V. Chain-of-Sight Implementation

In the COS problem, the different subtasks a robot can be allocated to are positions on the map shown in Fig. 1. Assuming discrete location values we could look at each pixel of the map as a potential location. However, since the location space $\mathcal{L}$ can be considerably large, depending on the dimensionality of the map, we regard only every m^{th} pixel as a valid location. It is now our goal to find the most probable location for each robot under the COS constraint, which is equivalent to approximating the marginals of each robot's hidden node x_i.

Recall the meaning of the local potential function $\phi(x_i, y_i)$. It expresses a robots own belief of being allocated to certain tasks, or in this case locations. We define this belief to be composed of the distance a robot would have to travel to get to a particular location, the desire not to move to another location at all because the robot is currently seeing the goal object, and the fact of a location being occupied (cf. Eq. 3). In Eq. 3 the term $D(x_i, y_i)$ denotes a "reciprocal" traveling distance, $G(x_i, y_i)$ represents goal attraction and $O(x_i, y_i)$ expresses occupancy of locations by other robots or obstacles. For the reciprocal traveling distance we map path lengths from the robot's current location l to any location in $\mathcal{L}$ linearly such that we obtain 1 for the shortest path and almost 0 for the longest path. Goal attraction is achieved by placing a Gaussian bell curve around the robot's current location if the goal object has been spotted. If no goal is spotted, $G(x_i, y_i)$ is uniform across all possible locations. The occupancy function is defined as

$$O(x_i, y_i) = M(x_i) \prod_{k \in N(i)} 1 - b_i(x_k) \tag{4}$$

where $M(x_i)$ returns 0 if a particular location is occupied by an obstacle, such as a wall or an object sensed through a range finder, and 1 if the location is not occupied. The second term of Eq. (4) allows for the incorporation of other robot's intentions. $b_i(x_k)$ encodes belief distributions about where other robots will probably be allocated. By taking this information into account robot i can avoid to allocate itself to locations which are already assigned to other robots. If for some reason the intentions of other robots are not available, $b_i(x_k)$ can be regarded as uniform across all locations.

In order to define the potential function $\psi_{i,j}(x_i, x_j)$, we address where robot i would like to see another robot j assuming robot i is at a particular location. Thus, for each possible location l that x_i can take on the potential function $\psi_{i,j}(x_i, x_j)$ returns a distribution that is 1) highly peaked around the location that is visible to l and 2) at the same time furthest away from l toward the starting location. We assume distances between locations and the starting point can be calculated using Dijkstra's shortest path algorithm. In order to determine visibility between l and all other locations we construct a visibility polygon around l and test if a particular location lies inside this polygon. Notice that once the location space and terrain map is known, $\psi_{i,j}(x_i, x_j)$ can be completely precomputed.

Each robot is constantly processing incoming messages and sending out messages to robots in its neighborhood. Processing a message involves the recomputation of a robots own belief and an allocation to the position that has the highest probability, that is, where the robots own belief is highest. When sending a message, a robot chooses the recipient at random from its neighborhood.

A. *Player Specifics*

We implemented MRBP as explained above for the Player framework [4] and physically simulated the behavior of five robots on the map shown in Fig. 1 using Gazebo. Each robot was controlled by its own process and communicated with other robots through UDP/IP. Once allocated to a location, a robot uses Player's wavefront planning proxy to navigate from its current position to the desired location. Local collision avoidance is performed through Player's vector field histogram proxy using laser range data.

VI. Results and Discussion

A. *Simulation Results*

We conducted a total of 50 COS trials in simulation with five robots, achieving a successful allocation in all of these runs. The 50 trials consisted of test runs on five different configurations of start, goal and robot locations, resulting in 10 trials per configuration. The start and goal locations as well as the position of the robot facing the goal object were chosen manually whereas the remaining four robots were positioned randomly. Depending on the complexity of the problem, that is how many robots were needed to construct a minimal COS, more messages had to be propagated to obtain a valid arrangement. The number of total messages transmitted before a correct arrangement arose ranged from 10 messages, in situations where only two robots were required for the COS, to 80 messages for complicated arrangements. The number of messages needed also varied over multiple runs on the same scenario, an indication for the randomness underlying the message recipient selection. The allocations typically changed to some extent as the robots moved toward their assigned locations, most probably due to the varying traveling distance component in the local evidence function. Once the group found a stable configuration the allocations rarely changed in a negative way, and when they did, they reconverged. Fig. 4 graphically shows the decomposition of a robot's belief during a test run.

In 10 of these 50 trials, a strategically important robot was purposely shut down to test dynamic recovery abilities of the system. When such faults occurred, the robots properly adapted their beliefs, which resulted in reallocations that "filled" the gaps. Note that for obvious reasons the group can only recover from such a failure if there are enough robots left to physically span a chain. In cases where there a more robots in a group than are needed to span a minimal COS, the "redundant" robots do not impede normal execution and allocate themselves along the chain-of-sight path.

B. *Open Issues*

Specifically to the COS problem there are the open issues of how to handle multiple goal objects and how to keep multiple robot's from being attracted to the same goal object. Also, it has not been discussed how to transition from the search phase in which robots spread out to the chain formation phase.

A general issue with MRBP is that there is no way for a robot to find out to what extent other robots followed its advice and allocated themselves according to its messages; at least not through MRBP as described here. Another issue relates to the problem of multiple robots being allocated to the same task. If multiple robots receive exactly the same message from a particular robot it can happen that several of those robots allocate themselves to the same task. In case of COS this means that a group of robots clusters around a specific location. As mentioned, we counteracted this behavior by

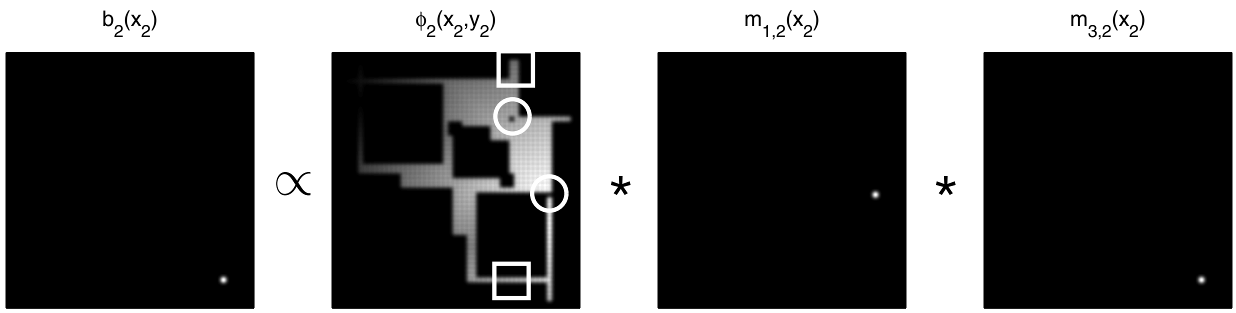

Fig. 4. Visualization of robot number 2's belief and its decomposition into local evidence and incoming messages during a test run of the COS implementation with three robots. Black pixels indicate low probability whereas white pixels indicate high probability. The goal and start location are marked by the upper and lower white square, respectively. The two white circles highlight the locations of robot 1 and 3 which are reflected in the local evidence in terms of occupied locations (hence the dark spots). The upper circle depicts robot 1 and the lower circle robot 3. The gradient from black to white in the local evidence is due to the traveling distance component. The message received from robot 1 indicates that robot 1 advices the robots in its neighborhood to allocate themselves to a location in the middle-right part of the map. Robot 2 does not "follow" this advice because that particular spot is already occupied by robot 3 (cf. the lower white circle in the local evidence). Robot 3's advice to its neighbors is to allocate themselves to a location in the lower-right corner. Since this spot is not occupied it yields the highest probability in robot 2's belief distribution.

including an occupancy term into the local evidence function. Another interesting solution to this problem could consist of extending the compatibility function $\psi_{i,j}(x_i, x_j)$ such that it is specific to the relationship between robot i and j. That is, robot i does not send the same message to all its neighbors but tailors recipient-specific messages. That way, robot i could focus its advice to a particular neighbor and thus avoid that multiple robots might follow a general advice.

COS is only one particular application of MRBP and so far no experiments have been performed with groups of real robots. We consider multi-robot belief propagation as a promising novel approach to distributed task allocation and note that it is still in its early, experimental stages. Future comparisons with existing MRTA algorithms and implementations for other allocation scenarios shall yield more insights into the robustness and applicability of MRBP.

VII. Conclusion

We have presented Multi-Robot Belief Propagation as a means to perform distributed multi-robot task allocation. We combine distributed statistical inference with theories of human development to address the task allocation problem. Through MRBP, collaborative robot group behavior results that accounts for both individual decision making and collective goal achievement. We demonstrated MRBP for an example chain-of-sight problem, whose generalizations we believe will be beneficial for a wide variety of multi-robot applications.

References

[1] T. Balch and R. C. Arkin. Communication in reactive multiagent robotic systems. *Autonomous Robots*, 1(1):27–52, 1994.

[2] M. E. Bratman. *Intentions, Plans, and Practical Reason*. Harvard, Cambridge, MA, USA, 1987.

[3] C. Breazeal, A. Brooks, J. Gray, G. Hoffman, C. Kidd, H. Lee, J. Lieberman, A. Lockerd, and D. Chilongo. Tutelage and collaboration for humanoid robots. *International Journal of Humanoid Robotics*, 1(2):315–348, June 2004.

[4] T. Collett, B. MacDonald, and B. Gerkey. Player 2.0: Toward a Practical Robot Programming Framework. *Proceedings of the Australasian Conference on Robotics and Automation (ACRA 2005)*, 2005.

[5] M. Dias and A. Stentz. Opportunistic optimization for market-based multirobot control. In *Proc. of the 2002 Int. Conf. on Intel. Rob. and Sys.*, volume 3, pages 2714–2720, September 2002.

[6] S. Geman and D. Geman. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6(6):721–741, 1984.

[7] B. P. Gerkey and M. J. Mataric. A formal analysis and taxonomy of task allocation in multi-robot systems. *I. J. Robotic Res.*, 23(9):939–954, 2004.

[8] B. Jung and G. S. Sukhatme. Tracking targets using multiple robots: The effect of environment occlusion. *Autonomous Robots*, 13(3):191–205, 2002.

[9] A. Leslie. ToMM, ToBy, and Agency: Core architecture and domain specificity. In L. Hirschfeld and S. Gelman, editors, *Mapping the mind: Domain specificity in cognition and culture*, pages 119–148. Cambridge University Press, 1994.

[10] A. Martinoli, A. J. Ijspeert, and L. M. Gambardella. A probabilistic model for understanding and comparing collective aggregation mechansims. In D. Floreano, J.-D. Nicoud, and F. Mondada, editors, *ECAL*, volume 1674 of *Lecture Notes in Computer Science*, pages 575–584. Springer, 1999.

[11] L. Parker. Alliance: An architecture for fault-tolerant multi-robot cooperation. *IEEE Transactions on Robotics and Automation*, 14(2):220–240, 1998.

[12] B. Scassellati. Theory of mind for a humanoid robot. *Auton. Robots*, 12(1):13–24, 2002.

[13] D. A. Shell and M. J. Mataric. Principled synthesis for large-scale systems: task sequencing. In *Proceedings of International Symposium on Distributed Autonomous Robotic Systems (DARS)*, July 2006.

[14] J. S. Yedidia, W. T. Freeman, and Y. Weiss. Understanding belief propagation and its generalizations. In G. Lakemeyer and B. Nebel, editors, *Exploring Artificial Intelligence in the New Millennium*, chapter 8, pages 239–269. Morgan Kaufmann, 2003.

[15] J. S. Yedidia, W. T. Freeman, and Y. Weiss. Constructing free-energy approximations and generalized belief propagation algorithms. *IEEE Trans. Inform. Theory*, 51(7):2282–2312, 2005.