ENABLING AUTONOMOUS SENSOR-SHARING FOR TIGHTLY-COUPLED COOPERATIVE TASKS

Lynne E. Parker, Maureen Chandra, and Fang Tang
Distributed Intelligence Laboratory, Department of Computer Science
The University of Tennessee, Knoxville, Tennessee USA
parker@cs.utk.edu, chandra@cs.utk.edu, ftang@cs.utk.edu

Abstract This paper presents a mechanism enabling robot team members to share sensor information to achieve tightly-coupled cooperative tasks. This approach, called ASyMTRe, is based on a distributed extension of schema theory that allows schema-based building blocks to be interconnected in many ways, regardless of whether they are on the same or different robots. The inputs and outputs of schema are labeled with an information type, inspired by the theory of information invariants. By enabling robots to autonomously configure their distributed schema connections based on the flow of information through the system, robot teams with different collective capabilities are able to generate significantly different cooperative control strategies for solving the same task. We demonstrate the ability of this approach to generate different cooperative control strategies in a proof-of-principle implementation on physical robots performing a simple transportation task.

Keywords: Sensor-sharing, heterogeneous teams, multi-robot coalitions

1. Introduction

In multi-robot systems, it is advantageous to be able to treat each sensory resource on the team as a resource available to any necessitous robot team member, rather than being exclusively owned by an individual robot. The ability to share sensory information, appropriately translated to another robot's perspective, can extend the task capabilities of a given multi-robot team. In practice, however, this is difficult to achieve because each sensory resource is in fact fixed on a particular robot, and provides information only from that robot's frame of reference. Typically, mechanisms for sharing distributed sensory information are developed in an application-specific manner. The human designer might pre-define roles or subtasks, together with a list of required capabilities needed to accomplish each role or subtask. The robot team members can then autonomously select actions using any of a number of common approaches to multi-robot task allocation (see (Gerkey and Mataric, 2004) for a comparison of various task allocation approaches), based upon their suitability for the role or subtask, as well as the current state of the multi-robot system. The shortcoming of this approach is that the designer has to consider in advance all of the possible combinations of robot capabilities that might be present on a multi-robot team performing a given task, and to design cooperative behaviors in light of this advance knowledge.

However, as described in (Parker, 2003), the specific robot capabilities present on a team can have a significant impact on the approach a human designer would choose for the team solution. The example given in (Parker, 2003) is that of deploying a mobile sensor network, in which cooperative solutions for the same task could involve potential-field-based dispersion, marsupial delivery, or assistive navigation, depending on the capabilities of the team members.

Our research is aimed at overcoming these challenges by designing flexible sensor-sharing mechanisms within robot behavior code that do not require task-specific, pre-defined cooperative control solutions, and that translate directly into executable code on the robot team members. Some related work in sensor-sharing has led to the development of application-specific solutions that allow a robot team member to serve as a remote viewer of the actions of other teammates, providing feedback on the task status to its teammates. In particular, this has been illustrated by several researchers in the multi-robot box pushing and material handling domain (Gerkey and Mataric, 2002, Adams et al., 1995, Spletzer et al., 2001, Donald et al., 1997), in which one or more robots push an object while a remote robot or camera provides a perspective of the task status from a stand-off position. Our work is aimed at generating these types of solutions automatically, to enable robot teams to coalesce into sensor-sharing strategies that are not pre-defined in advance.

Our approach, which we call ASyMTRe (Automated Synthesis of Multi-robot Task solutions through software Reconfiguration, pronounced "Asymmetry"), is based on a combination of schema theory (Arkin et al., 1993) and inspiration from the theory of information invariants (Donald et al., 1993). The basic building blocks of our approach are collections of *perceptual schemas*, *motor schemas*, and a simple new component we introduce, called *communication schemas*. These schemas are assumed to be supplied to the robots when they are brought together to form a team, and represent baseline capabilities of robot team members. The ASyMTRe system configures a solution by choosing from different ways of combining these building blocks into a teaming solution, preferring the solution with the highest utility. Different combinations of building blocks can yield very different types of cooperative solutions to the same task.

In a companion paper (Tang and Parker, 2005), we have described an automated reasoning system for generating solutions based on the schema building blocks. In this paper, we focus on illustrating a proof-of-principle task that shows how different interconnections of these schema building blocks can yield fundamentally different solution strategies for sensor-sharing in tightly-coupled tasks. Section 2 outlines our basic approach. Section 3 defines a simple proof of principle task that illustrates the ability to formulate significantly different teaming solutions based on the schema representation. Section 4 presents the physical robot results of this proof-of-principle task. We present concluding remarks and future work in Section 5.

2. Approach

Our ASyMTRe approach to sensor-sharing in tightly-coupled cooperative tasks includes a formalism that maps environmental, perceptual, and motor control schemas to the required flow of information through the multi-robot system, as well as an automated reasoning system that derives the highest-utility solution of schema configurations across robots. This approach enables robots to reason about how to solve a task based upon the fundamental information needed to accomplish the objectives. The fundamental information will be the same regardless of the way that heterogeneous team members may obtain or generate it. Thus, robots can collaborate to define different task strategies in terms of the required flow of information in the system. Each robot can know about its own sensing, effector, and behavior capabilities and can collaborate with others to find the right combination of actions that generate the required flow of information to solve the task. The effect is that the robot team members interconnect the appropriate schemas on each robot, and across robots, to form coalitions (Shehory, 1998) to solve a given task.

2.1 Formalism of Approach

We formalize the representation of the basic building blocks in the multi-robot system as follows:

- A class of *Information*, denoted $F = \{F_1, F_2, ...\}$.

- *Environmental Sensors*, denoted $ES = \{ES_1, ES_2, ...\}$. The input to ES_i is a specific physical sensor signal. The output is denoted as $O^{ES_i} \in F$.
- *Perceptual Schemas*, denoted $PS = \{PS_1, PS_2, ...\}$. Inputs to PS_i are denoted $I_k^{PS_i} \in F$. The perceptual schema inputs can come from either the outputs of communication schemas or environmental sensors. The output is denoted $O^{PS_i} \in F$.
- *Communication Schemas*, denoted $CS = \{CS_1,\ CS_2,\ ...\}$. The inputs to CS_i are denoted $I_k^{CS_i} \in F$. The inputs originate from the outputs of perceptual schemas or communication schemas. The output is denoted $O^{CS_i} \in F$.
- *Motor Schemas*, denoted $MS = \{MS_1, MS_2, ...\}$. The inputs to MS_i are denoted $I_k^{MS_i} \in F$, and come from the outputs of perceptual schemas or communication schemas. The output is denoted $O^{MS_i} \in F$, and is connected to the robot effector control process.
- A set of n robots, denoted $R = \{R_1, R_2, ..., R_n\}$. Each robot is described by the set of schemas available to that robot: $R_i = \{\boldsymbol{ES}^i$, $\boldsymbol{PS}^i$, $\boldsymbol{CS}^i$, $\boldsymbol{MS}^i\}$, where $\boldsymbol{ES}^i$ is the set of environmental sensors available to R_i, and $\boldsymbol{PS}^i$, $\boldsymbol{CS}^i$, $\boldsymbol{MS}^i$ are the sets of perceptual, communication, and motor schemas available to R_i, respectively.
- *Task* $= \{MS_1, MS_2, ...\}$, which is the set of motor schemas that must be activated to accomplish the task.

A valid configuration of schemas distributed across the robot team has all of the inputs and outputs of the schemas in T connected to appropriate sources, such that the following is true: $\forall_k \exists_i CONNECT(O^{S_i}, I_k^{S_j}) \Leftrightarrow O^{S_i} = I_k^{S_j}$, where S_i and S_j are types of schemas. This notation means that for all the inputs of S_j, there exists some S_i whose output is connected to one of the required inputs. In (Tang and Parker, 2005), we define quality metrics to enable the system to compare alternative solutions and select the highest-quality solution. Once the reasoning system has generated the recommended solution, each robot activates the required schema interconnections in software.

3. Proof-of-Principle Task Implementation

To show that it is possible to define basic schema building blocks to enable distributed sensor sharing and flexible solution approaches to a tightly-coupled cooperative task, we illustrate the approach in a very simple proof of principle task. This task, which we call the *transportation*

Table 1. Environmental Sensors (ES) and Robot Types for proof-of-principle task.

Environmental Sensors			Robot Types	
Name	Description	Info. Type	Name	Available Sensors
ES_1	Laser	*laserscanner*	R_1	Laser
ES_2	Camera	*ccd*	R_2	Camera
ES_3	DGPS	*dgps*	R_3	DGPS
			R_4	Laser and Camera
			R_5	Laser and DGPS
			R_6	Camera and DGPS
			R_7	Laser and Camera and DGPS
			R_8	—

task, requires each robot on the team to navigate to its pre-assigned, unique goal point. In order for a robot to reach its assigned goal, it needs to know its current position relative to its goal position so that it can move in the proper direction. In some cases, a robot may be able to sense its current position using its own sensors. In other cases, the robot may not have enough information to determine its current position. In the latter case, other more capable robots can help by sharing sensor information with the less capable robot.

As shown in Table 1, the environmental sensors available in this case study are a laser scanner, a camera, and Differential GPS. A robot can use a laser scanner with an environmental map to localize itself and calculate its current global position. A robot's camera can be used to detect the position of another robot relative to itself. The DGPS sensor can be used outdoors for localization and to detect the robot's current global position. Based upon these environmental sensors, there are eight possible combinations of robots, as shown in Table 1. In this paper, we focus on three types of robots – R_8: a robot that possesses no sensors; R_2: robot that possesses only a camera; and R_4: a robot that possesses a camera and a laser ranger scanner (but no DGPS).

For this task, we define five perceptual schemas, as shown in Table 2. PS_1 calculates a robot's current global position. With the sensors we have defined, this position could be determined either by using input data from a laser scanner combined with an environmental map, from DGPS, or from communication schemas supplying similar data. For an example of this latter case, a robot can calculate its current global position by knowing the global position of another robot, combined with its own position relative to the globally positioned robot. PS_2 outputs a robot's goal position, based on the task definition provided by the user.

Table 2. Perceptual and Communications Schemas for proof-of-principle task.

Perceptual Schemas		
Name	**Input Info. Type**	**Output Info. Type**
PS_1	*laserrange* OR *dgps* OR *curr-global-pos*(*self*) OR (*curr-rel-pos*($other_k$) AND *curr-global-pos*($other_k$))	*curr-global-pos*(*self*)
PS_2	—	*curr-global-goal*(*self*)
PS_3	(*curr-global-pos*(*self*) AND *curr-rel-pos*($other_k$))	*curr-global-pos*($other_k$)
PS_4	*laserrange* or *ccd*	*curr-rel-pos*($other_k$)
PS_5	*curr-global-pos*(*other*)	*curr-global-pos*(*other*)

Communication Schemas		
Name	**Input Info. Type**	**Output Info. Type**
CS_1	*curr-global-pos*(*self*)	*curr-global-pos*($other_k$)
CS_2	*curr-global-pos*($other_k$)	*curr-global-pos*(*self*)

PS_3 calculates the current global position of a remote robot based on two inputs – the position of the remote robot relative to itself and its own current global position. PS_4 calculates the position of another robot relative to itself. Based on the defined sensors, this calculation could be derived from either a laser scanner or a camera. PS_5 receives input from another robot's communication schema, CS_1, which communicates the current position of that other robot.

Communication schemas communicate data to another robot's perceptual schemas. As shown in Table 2, CS_1 communicates a robot's current global position to another robot, while CS_2 communicates the current global position of a remote robot that remote robot. Motor schemas send control signals to the robot's effectors to enable the robot to accomplish the assigned task. In this case study, we define only one motor schema, MS, which encodes a *go-to-goal* behavior.

The input information requirements of MS are *curr-global-pos*(*self*) and *curr-global-goal*(*self*). In this case, the motor schema's output is derived based on the robot's current position received from PS_1 and its goal position received from PS_2.

Figure 1 shows all the available schemas for this task and how they can be connected to each other, based on the information labeling. The solid-line arrows going into a schema represent an "OR" condition – it is sufficient for the schema to only have one of the specified inputs to produce output. The dashed-line arrows represent an "AND" condition, meaning that the schema requires all of the indicated inputs for it to calculate an output. For example, PS_1 can produce output with input(s)

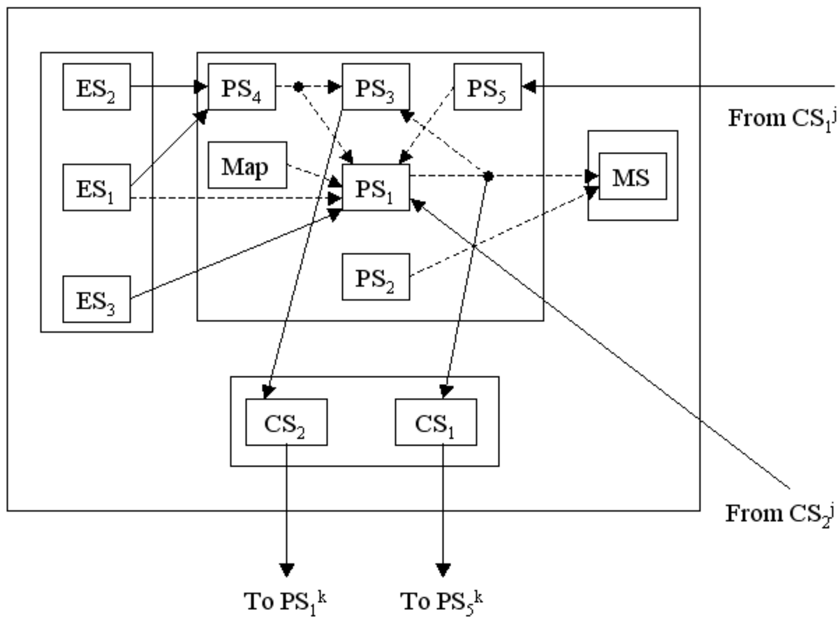

Figure 1. Illustration of connections between all available schemas.

from either ES_1 (combined with the environmental Map), ES_3, CS_2^j (R_j's CS_2), or (PS_4 and PS_5).

4. Physical Robot Experiments

These schema were implemented on two Pioneer robots equipped with a SICK laser range scanner and a Sony pan-tilt-zoom camera. Both robots also possessed a wireless ad hoc networking capability, enabling them to communicate with each other. Experiments were conducted in a known indoor environment using a map generated using an autonomous laser range mapping algorithm. Laser-based localization used a standard Monte-Carlo Localization technique. The code for the implementation of PS_4 makes use of prior work by (Parker et al., 2004) for performing vision-based sensing of the relative position of another robot. This approach makes use of a cylindrical marker designed to provide a unique robot ID, as well as relative position and orientation information suitable for a vision-based analysis. Using these two robots, three variations on sensor availability were tested to illustrate the ability of these building blocks to generate fundamentally different cooperative behaviors of the same task through sensor sharing. In these experiments, the de-

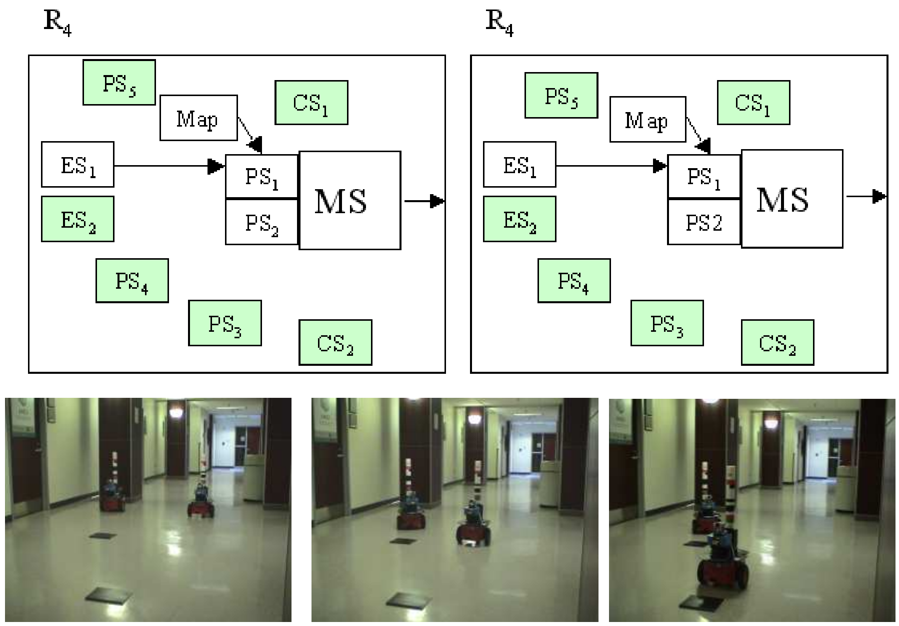

Figure 2. Results of Variation 1: Two robots of type R_4 performing the task without need for sensor-sharing or communication. Goals are black squares on the floor. Graphic shows schema interconnections (only white boxes activated).

sired interconnections of schemas were developed by hand; in subsequent work, we can now generate the required interconnections automatically through our ASyMTRe reasoning process (Tang and Parker, 2005).

Variation 1. The first variation is a baseline case in which both robots are of type R_4, meaning that they have full use of both their laser scanner and a camera. Each robot localizes itself using its laser scanner and map and reaches its own unique goals independently. This case is the most ideal solution but only works if the both robots possess laser scanners and maps. If one of the robots loses its laser scanner, this solution no longer works. Figure 2 shows the schema instantiated on the robots for this variation. PS_1 and PS_2 are connected to MS to supply the required inputs to the *go-to-goal* behavior. Also shown in Figure 2 are snapshots of the robots performing this instantiation of the schema. In this case, since both robots are fully capable, they move towards their goals independently without the need for any sensor sharing or communication.

Variation 2. The second variation involves a fully capable robot of type R_4, as well as a robot of type R_2 whose laser scanner is not available,

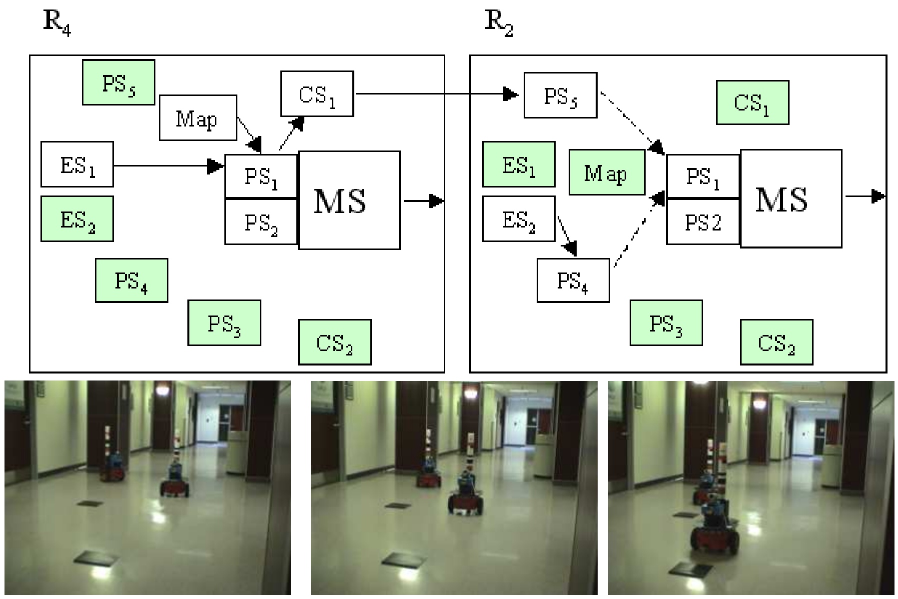

Figure 3. Variation 2: A robot of type R_4 and of type R_2 share sensory information to accomplish their task. Here, R_2 (on the left) turns toward R_4 to localize R_4 relative to itself. R_4 communicates its current global position to R_2, enabling it to determine its own global position, and thus move successfully to its goal position.

but still has use of its camera. As illustrated in Figure 3, Robot R_4 helps R_2 by communicating (via CS_1) its own current position, calculated by PS_1 using its laser scanner (ES_1) and environmental map. Robot R_2 receives this communication via PS_5 and then uses its camera (ES_2) to detect R_4's position relative to itself (via PS_4) and calculate its own current global position (using PS_1) based on R_4's relative position and R_4's communicated global position. Once both robots know their own current positions and goal positions, their motor schemas can calculate the motor control required to navigation to their goal points. Figure 3 also shows snapshots of the robots performing the Variation 2 instantiation of the schema. In this case, R_2 begins by searching for R_4 using its camera. At present, we have not yet implemented the constraints for automatically ensuring that the correct line of sight is maintained, so we use communication to synchronize the robots. Thus, when R_2 locates R_4, it communicates this fact to R_4. R_4 then is free to move towards its goal. If R_2 were to lose sight of R_4, it would communicate a message to R_4 to re-synchronize the relative sighting of R_4 by R_2. With this solution, the robots automatically achieve navigation assistance of a less capable robot by a more capable robot.

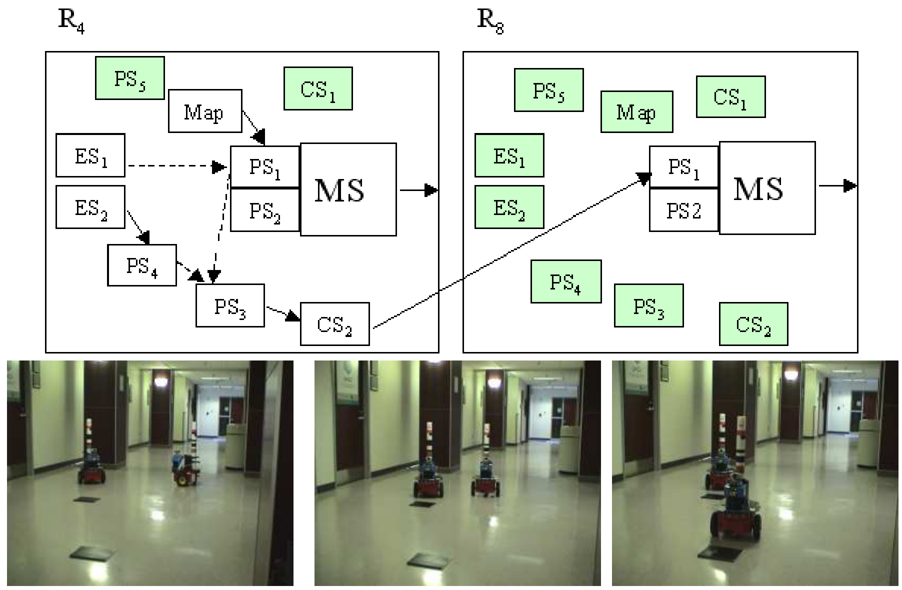

Figure 4. Variation 3: A robot of type R_4 helps a robot with no sensors (type R_8) by sharing sensory information so that both robots accomplish the objective. Note how R_4 (on the right) turns toward R_8 to obtain vision-based relative localization of R_8. R_4 then guides R_8 to its goal position. Once R_8 is at its goal location, R_4 then moves to its own goal position.

Variation 3. The third variation involves a sensorless robot of type R_8, which has access to neither its laser scanner nor camera. As illustrated in Figure 4, the fully-capable robot of type R_4 helps R_8 by communicating R_8's current global position. R_4 calculates R_8's current global position by first using its own laser (ES_1) and map to calculate its own current global position (PS_1). R_4 also uses its own camera (ES_2) to detect R_8's position relative to itself (using PS_4). Then, based on this relative position and its own current global position, R_4 calculates R_8's current global position (using PS_3) and communicates this to R_8 (via CS_2). Robot R_8 feeds its own global position information from R_4 directly to its motor schema. Since both of the robots know their own current and goal positions, each robot can calculate its motor controls for going to their goal positions. Figure 4 also shows snapshots of the robots performing the Variation 3 instantiation of the schema. With this solution, the robots automatically achieve navigation assistance of a sensorless robot by a more capable robot.

Analysis. In extensive experimentation, data on the time for task completion, communication cost, sensing cost, and success rate was collected as an average of 10 trials of each variation. Full details are available in (Chandra, 2004). We briefly describe here the success rate of each variation. In all three variations, robot R_4 was 100% successful in reaching its goal position. Thus, for Variation 1, since the robots are fully capable and do not rely on each other, the robots always succeeded in reaching their goal positions. In Variation 2, robot R_2 succeeded in reaching its goal 6 times out of 10, and in Variation 3, robot R_8 successfully reached its goal 9 times out of 10 tries. The failures of robots R_4 and R_8 in Variations 2 and 3 were caused by variable lighting conditions that led to a false calculation of the relative robot positions using the vision-based robot marker detection. However, even with these failures, these overall results are better than what would be possible without sensor sharing. In Variations 2 and 3, if the robots did not share their sensory resources, one of the robots would never reach its goal position, since it would not have enough information to determine its current position. Thus, our sensor sharing mechanism extends the ability of the robot team to accomplish tasks that otherwise could not have been achieved.

5. Conclusions and Future Work

In this paper, we have shown the feasibility of the ASyMTRe mechanism to achieve autonomous sensor-sharing of robot team members performing a tightly-coupled task. This approach is based on an extension to schema theory, which allows schemas distributed across multiple robots to be autonomously connected and executed at run-time to enable distributed sensor sharing. The inputs and outputs to schemas are labeled with unique information types, inspired by the theory of information invariants, enabling any schema connections with matching information types to be configured, regardless of the location of those schema or the manner in which the schema accomplishes its job. We have demonstrated, through a simple transportation task implemented on two physical robots, the ability of the schema-based methodology to generate very different cooperative control techniques for the same task based upon the available sensory capabilities of the robot team members. If robots do not have the ability to accomplish their objective, other team members can share their sensory information, translated appropriately to another robot's frame of reference. This approach provides a framework within which robots can generate the highest-quality team solution for a tightly-coupled task, and eliminates the need of the human designer to pre-design all alternative solution strategies. In continuing work, we are extending the formalism to impose motion constraints (such as line-of-sight) needed to ensure that robots can successfully share sensory data while they are in motion, generalizing the information labeling technique, and implementing this approach in more complex applications. In addition, we are developing a distributed reasoning approach that enables team members to autonomously generate the highest-quality configuration of schemas for solving the given task.

References

Adams, J. A., Bajcsy, R., Kosecka, J., Kumar, V., Mandelbaum, R., Mintz, M., Paul, R., Wang, C., Yamamoto, Y., and Yun, X. (1995). Cooperative material handling by human and robotic agents: Module development and system synthesis. In *Proc. of IEEE/RSJ International Conference on Intelligent Robots and Systems*.

Arkin, R. C., Balch, T., and Nitz, E. (1993). Communication of behavioral state in multi-agent retrieval tasks. In *Proceedings of the 1993 International Conference on Robotics and Automation*, pages 588–594.

Chandra, M. (2004). Software reconfigurability for heterogeneous robot cooperation. Master's thesis, Department of Computer Science, University of Tennessee.

Donald, B. R., Jennings, J., and Rus, D. (1993). Towards a theory of information invariants for cooperating autonomous mobile robots. In *Proceedings of the International Symposium of Robotics Research*, Hidden Valley, PA.

Donald, B. R., Jennings, J., and Rus, D. (1997). Information invariants for distributed manipulation. *International Journal of Robotics Research*, 16(5):673–702.

Gerkey, B. and Mataric, M. (2002). Pusher-watcher: An approach to fault-tolerant tightly-coupled robot cooperation. In *Proc. of IEEE International Conference on Robotics and Automation*, pages 464–469.

Gerkey, B. and Mataric, M. J. (2004). A formal analysis and taxonomy of task allocation in multi-robot systems. *Int. J. of Robotics Research*, 23(9):939–954.

Parker, L. E. (2003). The effect of heterogeneity in teams of 100+ mobile robots. In Schultz, A., Parker, L. E., and Schneider, F., editors, *Multi-Robot Systems Volume II: From Swarms to Intelligent Automata*. Kluwer.

Parker, L. E., Kannan, B., Tang, F., and Bailey, M. (2004). Tightly-coupled navigation assistance in heterogeneous multi-robot teams. In *Proceedings of IEEE International Conference on Intelligent Robots and Systems*.

Shehory, O. (1998). Methods for task allocation via agent coalition formation. *Artificial Intelligence*, 101(1-2):165–200.

Spletzer, J., Das, A., Fierro, R., Tayler, C., Kumar, V., and Ostrowski, J. (2001). Cooperative localization and control for multi-robot manipulation. In *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, Hawaii.

Tang, F. and Parker, L. E. (2005). ASyMTRe: Automated synthesis of multi-robot task solutions through software reconfiguration. *Submitted*.