

Murdoch: Publish/Subscribe Task Allocation for Heterogeneous Agents

Brian P. Gerkey and Maja J Matarić
Computer Science Department
University of Southern California
Los Angeles, CA 90089-0781
{bgerkey,mataric}@usc.edu

1. INTRODUCTION

In this paper, we describe a novel approach to the problem of dynamic task allocation among groups of heterogeneous agents. Specifically, we advocate the use of *publish/subscribe* messaging, a well-researched ([3]) and commercially proven ([4]) message brokering paradigm that is readily applicable to distributed control. We present MURDOCH, an implemented publish-subscribe system, and explain how it can facilitate multi-robot coordination.

The system allows the user a high-level interface for posing tasks to a group of autonomous homogeneous or heterogeneous robots. Rather than assigning a task to an individual or a group, the user simply poses the task to the system as a whole (with no knowledge of individual agents). To take full advantage of the publish/subscribe framework, we apply MURDOCH at the *behavior level* of the control system, rather than at the level of the robots. Behavior-based control has been shown to be a particularly effective substrate for distributed control, due to its scaling properties [2], [1].

2. PUBLISH/SUBSCRIBE MESSAGING

At the heart of MURDOCH is an implementation of publish/subscribe messaging, which in turn depends on subject-based addressing. *Subject-based addressing* is an addressing scheme in which individual messages are addressed by content rather than destination. *Publish/subscribe messaging* is a messaging paradigm that uses *subject-based addressing* to divide a network into a loosely-coupled association of anonymous data producers and consumers. A data producer simply tags a message with a subject (or set of subjects) and “publishes” it onto the network; any data consumers who have “subscribed” to that subject (or set of subjects) will automatically receive the message. The goal of publish/subscribe messaging is to enable a loosely-coupled distributed system in which the data producers have no knowledge of the data consumers and vice versa.

3. SUBJECT NAMESPACE

In our system, each robot subscribes to a set of subjects which represent its “capabilities”. A mobile robot might subscribe to subjects such as *sonar*, *speech*, *camera*, *compass*, and *mobile*; a desktop PC, on the other hand, might subscribe to subjects such as *compute-server* and *mass-storage*. Thus, to reach all robots who have sonar sensors and cameras, one can publish a message to the subject (*sonar camera*); only those machines with the specified capabilities will receive the message. In addition to those subjects which represent concrete robot capabilities, subjects can be more abstract representations of state information. For example, when a team of robots is engaged a cooperative box-pushing task, a robot’s current perception of the box is extremely useful; thus domain-specific subjects such as *see-box* and *contacting-box* will be introduced and used to coordinate the robots’ behavior throughout the progression of the task. These subscriptions can and do change over time. For example, with regard to our box-pushing task, which is a truly dynamic domain, individual sub-tasks will be periodically generated and published to subjects such as *contacting-box*; an agent’s state with respect to this subscription will be constantly changing throughout the task, and sub-task assignments will reflect this.

4. TASK ALLOCATION

4.1 Resource Requirements

In MURDOCH, task allocation is performed by matching the set of resources needed to perform the task with agents who are capable of achieving. As described in Section 3, an agent’s capabilities at any point in time are described by the subjects to which it subscribes. So, in order to match a given task to capable agents, messages regarding the task should be tagged with a subject that describes the necessary resources; specifically, those messages should be published to the subject that is the conjunction of the required resources. We consider resource requirements specified in this way (i.e. in the subject of a message) to be “hard”, in that they are not negotiable; if an agent does not subscribe to one of the subjects to which the message is addressed, then that agent must not have some necessary resource for the task at hand and so should not receive correspondence concerning the task.

Of course, more than one agent may be capable of performing a task and so there must be a method for selecting among a group of willing participants. We introduce metric func-

tions, or simply *metrics* for this purpose. Along with a task description, a message concerning a task to be performed will also include a set of metrics which will be evaluated on each respondent, assigning it a “score”, so that the system may (in a distributed fashion) select the agent best suited. As with task decomposition, selection of hard and soft requirements is the responsibility of the user, although it should be possible to extract at least hard requirements by examining the usage of primitives.

Metrics can take many forms, with the restriction that each one, when evaluated in the context of a specific agent, should return a scalar value (integral or real) representing that agent’s fitness for the task (in the current implementation, smaller values represent better fitness). A metric is usually defined as some function of the agent’s current state, although in general, a metric could perform any arbitrary computation, including inter-agent communication. Multiple metrics can be specified for one task and the order in which they are specified defines an implicit relative priority. Since it is our desire to derive a single scalar score for each agent on a given task, we must somehow combine the outputs from multiple metrics. At the moment, we use a fixed-weight linear combination that reflects relative the relative priorities (i.e. we compute a weighted sum of metric values, with greatest weight lent to the metric of highest priority). Other combination methods are certainly possible.

4.2 Negotiation

We now briefly describe the distributed negotiation protocol by which task allocation actually occurs within MURDOCH. The first step in using our system is to phrase the task to be performed in terms of task components, or sub-tasks, with each described in terms of behavioral primitives and hard and soft resource requirements. It should be noted that MURDOCH is a system for task *allocation*, not for task *decomposition*. Task decomposition is not currently autonomously handled by our system. Instead, we assume that the user performs this step and gives the resultant task description to MURDOCH for assignment. However, this process could certainly be automated through the use of, for example, a top-down planner.

Given such a description of the task, MURDOCH iterates through the list of sub-tasks, publishing for each a message announcing that it is available. Each *announcement* message is published to the subject that represents its hard resource requirements and includes: the sub-task description (primitive + parameterization), the relevant metrics, and a new, unique, subject (called a *sub-task channel*) for later communication concerning the sub-task.

Upon receipt of such an announcement message, an eligible agent first unsubscribes from the subjects to which it is addressed. The reason for the unsubscription is so that those resources are reserved in case the agent wins this sub-task in the end; reserving them by stopping listening to those subjects guarantees that the agent will not overextend itself by agreeing to do more than it is able. The agent next subscribes to the sub-task channel and then scores itself with respect to the sub-task by evaluating each of the supplied metric functions, and computing the weighted sum of the results. The agent publishes this value in a *score* message to

the sub-task channel, and waits for scores from other agents. All eligible agents perform these same steps, resulting in all their scores being broadcast to each other.

Since every agent who received the announcement must be eligible, and since every agent received and evaluated the same metric functions (we assume that the agents honestly report their own scores), we can conclude that the agent who published the lowest score (lower values are better) is best suited for performing the sub-task at hand. We employ a winner-take-all approach, and so once all the scores are in, every agent knows who the winner is. The winning agent proceeds to perform the sub-task while the other agents re-subscribe to the subjects corresponding to the temporarily reserved resources since they are once again available.

5. SUMMARY AND ONGOING WORK

In this paper, we have proposed the idea of applying a well-known idea from distributed system, namely publish/subscribe messaging, to the problem of task allocation for multiple heterogeneous agents. We have shown how this communication model is well-suited to such a control problem. Currently, we are working on a comprehensive demonstration of the utility of MURDOCH; our experimental domain is a cooperative planar object manipulation problem involving three Pioneer 2 mobile robots and making use of dynamic creation and assignment of tasks.

As for MURDOCH, we have specific plans for the future of MURDOCH, including the introduction of support for exclusive and non-exclusive reservation of resources, functionality for intelligent error detection (simple to do with “heartbeat” messages) and recovery (more difficult), and the ability to send messages (and metrics) with executable content.

Further information about this and related ongoing work is available at:

<http://www-robotics.usc.edu/~agents/pub-sub.html>

6. ACKNOWLEDGEMENTS

The research reported here was conducted at the Interaction Lab, part of the Robotics Research Lab at the University of Southern California Computer Science Department. The work is supported in part by the Office of Naval Research Grant N00014-95-1-0759 and in part by the DARPA Contract DAAE07-98-C-L028 and DARPA Grant DABT63-99-1-0015.

7. REFERENCES

- [1] R. C. Arkin. *Behavior-Based Robotics*. MIT Press, Cambridge, MA, 1998.
- [2] M. J. Matarić. Behavior-based control: Examples from navigation, learning, and group behavior. *Journal of Experimental and Theoretical Artificial Intelligence*, 9(2-3):323-336, 1997.
- [3] R. Strom et al. Gryphon: An Information Flow Based Approach to Message Brokering. In *Proc. of the International Symposium on Software Reliability Engineering*, Paderborn, Germany, Nov. 1998.
- [4] TIBCO Software, Inc. *TIB/Rendezvous Concepts*, Aug. 1997.