

Scaling High Level Interactions Between Humans and Robots¹

Ashley D. Tews, Maja J. Matarić and Gaurav S. Sukhatme

Center for Robotic and Embedded Systems, USC Robotics Research Lab,
Computer Science Department, University of Southern California,
Los Angeles, USA
{atews|mataric|gaurav}@robotics.usc.edu

Abstract

Much of the current research in human-robot interaction is concerned with single systems and single or few users. These systems and their interfaces are generally tightly-coupled and well-defined. For large-scale human-robot interactions, it may not be possible to know the systems in advance of designing the interface for potential human interaction. This presents a difficult goal for allowing multiple users to interact with many possibly unknown systems. In this paper, we present an interaction infrastructure aligned with providing this interface. It operates in two phases accommodating both many-to-many interactions and allowing for generalized and more personal one-to-one interaction between users and robotic systems. Our previous research has demonstrated the infrastructure to scale to a large number of users and several systems in simulation. The experiments in this paper substantiate these results in a smaller-scale real robotic environment.

Introduction

When considering large-scale human-robot interactions (in the order of hundreds of entities), it cannot be assumed that all robotic systems' interaction requirements will be known in advance. Therefore, any interaction mechanism used for this purpose must be general enough to accommodate these unknown factors. Typical human-system and human-robot interaction research relies on the constraints that the systems *are* known in advance. The interaction mechanism, or *infrastructure* presented in this paper relaxes these constraints and provides a general approach towards allowing many-to-many interactions between robotic systems and users as well as more personal one-to-one. To consider such large-scale interactions, we assume the systems are either semi- or fully autonomous.

Many issues influence the design of a multi-entity interaction mechanism, including the types of interaction to be catered for, the anticipated number of users and robots, and the heterogeneity of the interaction. In this section, the interaction discussion applies to both a single and group of robots which constitute a single entity for interaction purposes.

The type of interaction refers to the level of control the user has over the robot during interaction. For high-level control, the user has few control commands over the robot; these being high level requests for complex services the robot can provide. Supervisory control is a typical example where the user issues simple commands such as 'clean up room 214' and observes the robot carrying out the task. This type of interaction assumes the robot is partially or completely autonomous. Since the user is not bound to the robot during task execution, the interaction can be considered to be loosely-coupled.

More tightly-coupled interaction is evident at the low-level end of the spectrum. With this type of interaction, the user has a more intimate relationship with the robot by being able to access more of its lower-level control functions. Tele-operation (Murphy and Rogers 1996) of a mobile robot is a fundamental example where the user can control many of the robot's actions.

The number of entities for which the interaction mechanism is intended can range from a single robot and single user (one-to-one), to multiple robots and multiple users (many-to-many). Extending the numbers on either side of the interaction has a large impact on the mechanism design, particularly for tightly-coupled interaction where the dedicated interaction resources (such as control and communication) between a robot and user would have to be shared. Another issue that emerges from multi-entity interaction is its homogeneity. Homogenous interaction implies that the same level of control is allowed over each robot irrespective of their architecture. Heterogeneous interaction, in contrast implies that different levels of control are allowed.

Each of these issues has an impact on the complexity of a multi-entity interaction mechanism. We define multi-user, multi-robot interaction as *the ability for multiple users to interact with multiple robots simultaneously*. This definition requires addressing each of the issues described above. From analyzing these issues, we have determined that an interaction mechanism must be simple for scalability purposes but provide adequate functionality to contend with the possible types of interaction between

users and robots, and also the heterogeneity introduced when addressing scalability. Furthermore, the mechanism should provide extended functionality to take advantage of the opportunities of collaborating with multiple systems.

We have designed an interaction infrastructure with these goals in mind. It consists of a two-phased approach that allows the exposure of multiple robot systems to multiple users in the first phase, and individual user interaction with systems in the second. It also provides the extended functionality of task queuing and linking. An added benefit is that systems require little modification to be able to use the infrastructure.

Our previous work demonstrates the infrastructure's operation for tens of simultaneous users and services offered by robotic and embedded systems, and the advantages of queuing tasks in (Tews, Sukhatme, and Mataric 2002). The infrastructure's task-linking ability is demonstrated in (Tews et al. 2002). The research presented in this paper substantiates these experiments in the real robot domain with a smaller number of entities. Both scalability and task-linking experiments are reproduced and demonstrate the infrastructure as being a realistic scalable interaction mechanism.

Related Work

In this section, we classify human-system architectures in terms of their scalability and assume that a system consists of a set of nodes (e.g. robots) that provide a set of services to users. These classifications are shown in Table 1 and discussed next as single (non-shaded table entries) and multi-user architectures (shaded table entries).

Table 1. User/system scalability classification

	Users	Systems
1	One	One
2	One	Multiple
3	Multiple	One
4	Multiple	Multiple

Single-user System Architectures

There are many examples of research relating to the entries in the table above. The issues concerning scalability are highlighted when examining them. As mentioned in the previous section, an important issue that affects the scalability of an approach is how *tightly-coupled* the users are with the system during interaction. This concerns both access to and control of system resources. Single-user systems can allow the user full access to its resources. Adding more users means these resources have to be shared which can be detrimental or prohibitive for some types of interaction, particularly those that require a high and dedicated communication bandwidth.

Multi-robot research provides examples where robots can be controlled individually or as a group by a single user (Ali and Arkin 1994; Arkin, Collins, and Endo 1999; Jones and Snyder 2001; Rybski et al. 2001). In each of these systems, the robots have a certain degree of

autonomy and the user has either high-level group control, or low-level control of an individual robot. This indicates that the higher the number of entities that an interaction mechanism has to deal with, the looser their coupling should be.

Multi-user System Architectures

Multi-user and/or multi-system interaction is forced to address issues of scalability, resource allocation, and distributed control. A key concern is how to manage the information and communication pathways for multiple entities. Common approaches involve using multi-agent architectures and the Internet.

Multi-agent architectures such as JINI, CHI (Goldsmith, Spires, and Phillips 1998; Phillips, Goldsmith, and Spires 1998) and Icraftor (Ponnekanti et al. 2001) offer services to users with suitable access to the systems they serve. An example of their application is for interactive workspaces (Fox et al. 2000; Lee and Hashimoto 2000) which consist of embedded devices in dedicated environments that provide services to users. The users typically must physically be present in the environment to access the services.

Multi-user, multi-robot architectures are demonstrated by (Elhajj et al. 2000; Surmann and Theißinger 1999). These systems use the Internet as a communication medium for commanding robots from remote locations.

Many of the examples presented in this section are application- or domain-dependent. If they allow multiple users, it is typically in pre-determined, small numbers. The interaction infrastructure we present in this paper operates beyond these limits by placing little constraint on the number of users, systems and heterogeneity of the possible interactions.

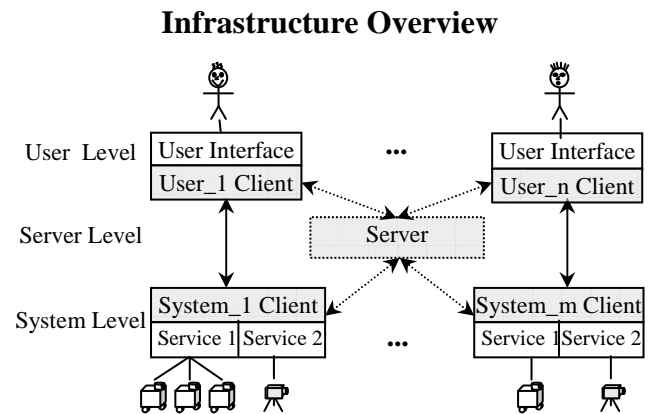


Figure 1. An overview of the infrastructure for connecting multiple users to multiple system services. Phase one is shown by broken lines and potential phase two interactions by solid lines.

Our infrastructure is based on a server/client architecture as shown in Figure 1. Each system is typically a robot, robot group, or embedded system (e.g., a security camera network). Each entity (system or user) has a client that interfaces with the server and manages low-level signal

and message passing as shown in Figure 2. It is assumed that the server's contact information is known to all systems and users wishing to connect to the infrastructure. The server is the centralized hub of operations that manages the allocation of systems' services and users' requests. The data flow throughout the infrastructure is by UDP communication.

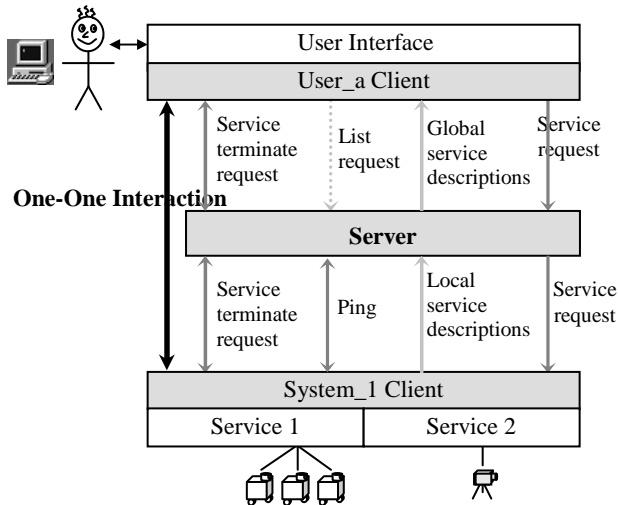


Figure 2. Signal and data flow through the infrastructure between a single system and user. The dashed lines are signals and the solid lines are messages.

There are two phases to the infrastructure's operation: many-to-many and one-to-one interaction. The first phase establishes indirect contact between users and systems via the server. Upon connection, each system provides a list of its offered capabilities – termed 'services' to the server and whether a task queue should be set up for it. The service list contains the description and type of each service with types being either 'data' or 'command'. Data services are typically provided by autonomous or stand-alone system nodes and inform the user of their task status during operation. Command services require the user to have a more active role by providing information pertaining to their tasks and are typically semi-autonomous.

The server maintains a database of all systems, their service lists, and contact information. Systems are routinely contacted to ensure they are still serviceable. Any that do not respond in a reasonable amount of time are removed from the server and notified. They can easily reconnect if they received this message and are still operational. This disconnection/reconnection does not affect existing interactions between the system and users.

When users connect to the infrastructure, they are provided with the current global list of all services in the server's database. The list is presented to them as a menu of options and includes the extended options of refreshing the list, service-linking, and exiting; each of which will be discussed next.

Users need to periodically refresh their service lists since systems and services will be added and removed

constantly from the server's database. Hence, the users will only have a 'snapshot' of the database at any point in time.

Linking services allows the user to combine two services to perform a higher-level function. The services consist of a command and data type. The data service provides relevant information to the command service. They must be semantically compatible with each other to allow service linking. This is currently the user's responsibility.

The second phase of operation occurs when the user selects a service from the menu. The selection is processed by the server to determine whether the service is available or not. The service is available if the user is the only one on an established task queue, or there is no queue. If the service has been removed from the server or is unavailable, the user is notified. Otherwise, the request is forwarded to the applicable system. A system- and service-dependent level (coupling) of interaction is then established directly between the user and the system. The communication bandwidth between interacting entities may be reduced if the user selects multiple services, or the system is managing multiple requests. Once a service has been selected, the system may send a signal to the server to remove it from the global service list. This mainly applies to systems and services that can only manage a single request. Interaction continues between the service and user until either terminates it. The user may select services at any time, resulting in one-to-many interaction. More details on the infrastructure's operation are provided in (Tews, Sukhatme, and Matarić 2002).

Experiments

In previous experiments, we demonstrated the infrastructure's scalability (Tews, Sukhatme, and Matarić 2002) and service-linking abilities (Tews et al. 2002). The scalability experiment demonstrated hundreds of users using the infrastructure to access 80 services offered by robotic and embedded systems. The service-linking experiment demonstrated the advantages of connecting services through the infrastructure where a robotic tracking system's service provided target location information to a robot offering a delivery service. These experiments were carried out using the Stage simulator and Player devices (Gerkey et al. 2001). Stage has been demonstrated on many projects to provide a valid representation of the real operation of Pioneer robots. Player provides a suite of software clients that represent the handlers for the devices on the robots such as sonar, SICK laser range-finders and color cameras. The combination of Player and Stage allows programs developed and tested in Stage to transfer readily to the real robots.

The experiments in this section substantiate the previous simulation results in a real environment on a smaller scale. Figure 3 shows the environment with the robots and tracking system coverage. Four Pioneer 2 DX robots are used to provide tracking and delivery systems. Each robot is equipped with an on-board sonar ring, SICK laser range-finder and Sony PTZ camera. The tracking

robots use the cameras and ACTS vision system to locate and identify colored beacons on targets. The laser range-finders are used with the robot's odometry for laser beacon-based localization and navigation. The sonar is used for obstacle avoidance.




Figure 3. The Stage simulated view of the office floor environment showing the tracking robots (T1-3) and delivery robot (D1) in their initial positions. The gray areas indicate the range of each tracking robot.

The first system is a multi-robot, multi-target tracking system (Jung and Sukhatme 2002), configured to detect three different targets identified by colored beacons. Hence, three services are generated and provide target location information. For the experiments, the system consists of three robots configured to be either stationary or dynamic trackers. Using the infrastructure, the system client can manage up to 10 requests simultaneously.

The second system offers only one service and consists of a delivery robot that navigates to goal locations provided by the requesting user. It provides the task status details of proximity to the goal location and the robot's current pose. A new goal location can be entered at any time. The delivery robot has a 'home' location in one of the offices where it returns to whenever it is not tasked. For long duration experiments, this location could contain a docking station for recharging.

The two systems share the same world map and coordinate reference frame so location information is common between them. This map was used in the earlier simulation experiments. No calibration had to be carried out for real robot localization, which is a testament to the accuracy of the map used in Stage and the effectiveness of using Stage for real robot system development.

Experiment 1 – Scalability Analysis

The first experiment involves many users requesting tasks from the systems. The three tracking robots are configured to act as stationary sensors. The area of the tracking system's coverage is shown by the gray areas in Figure 3. The target to be tracked is the delivery robot which is tasked through the infrastructure to traverse the corridors within the tracking system's coverage.

The systems connect to the infrastructure, register their services and wait for requests. The users connect to the infrastructure and receive the global list of services,

consisting of the three tracking services for the differently colored targets and one delivery service. The users select whichever service they desire and are connected to the applicable system. They then have a one-to-one interaction for each system they are connected to. Additionally, they retain access to the global service list for the many-to-many interaction.




Figure 4. Server response times for Experiment 1. The number of users connected at any point in time is represented by the solid line and the right y axis. The server's response times for the requests is shown by the symbols and correlate to the left y axis.

The experiment ran for approximately 50 minutes with 10 users connecting to the infrastructure and periodically requesting and terminating services. Figure 4 shows the server's performance during the experiment. The data shown represents the time for the server to process service request, list request, and service termination messages. The instantaneous number of online users is also shown. As can be seen, the number of users has little effect on the response times. The main reason for this is the absence of user maintenance, since the infrastructure's operations are event-based on the user's side. Therefore, connected users require no maintenance when there is no activity. On the system side, a ping signal is sent from the server every 20 seconds to ensure all systems connected to it are still serviceable. A system has 10 seconds to respond before it is removed from the server.

Table 2. Statistics from the user's perspective

	Service List Requests	Service Requests
Number of Requests	30	19
Maximum Response Time	170 ms	64 s
Median Response Time	43 ms	691 ms
Minimum Response Time	5 ms	38 ms

While the server's response times to requests provide an indication of its efficiency, the response times from the users' perspective provides a better indication to the efficiency of the infrastructure. In addition to the time taken by the server to process the requests, the

communication delays between infrastructure layers, system client processing times and UDP packet queue processing can also contribute to the response times. Table 2 lists the statistics pertaining to requests for services and service lists. These constitute the two most time-intensive messages used in the infrastructure. The times are measured from when the user made the request to when a response was received which is either from the server if the request is refused or they are on a queue, or from the system to commence interaction.

Requests for the global service list from the server are generally carried out quickly since the transaction involves only the user clients and the server. Requests for services take longer since they are processed by the server and the system clients. Apart from the transition times, it may take some time for the system client to respond to the user's request depending on its ongoing activities. For the tracking system, one robot was responsible for managing all users' requests in addition to its tracking duties. As a consequence of these duties and the limited processing capability and memory on the robot, the response times sometimes took several seconds.

The statistics for the service requests are listed in Table 3. The users selected 19 tasks, 14 of which were granted. Three tracking service requests were refused due to the maximum quota of 10 simultaneous requests being exceeded². The system client notified the users in these cases. Two request failures occurred due to the robot being overburdened with too many requests to manage simultaneously. Although the maximum number of requests that the system client manages was set to 10, the maximum number the robot's internal system could reasonably sustain while carrying out its part as a tracker was eight. During peak activity, eight users were being managed by the tracking system for more than 20 minutes, providing updates at 1 Hz.

Table 3. Service statistics showing the number of tasks selected and service running times

	Delivery Service	Tracking Service
Number of Granted Tasks	3/3	11/16
Maximum Time	1754 s	2057 s
Median Time	341 s	1478 s
Minimum Time	185 s	2 s

Experiment 2 – Service Linking

The second experiment demonstrates the ability of the infrastructure to link heterogeneous services to provide a higher-level function to the user. It is similar to our previous service-linking experiment (Tews et al. 2002) except it is carried out in a real building environment rather than simulated open space.

² Even though there were only 10 users, a user could request more than one task from a system if the system allows it.

In this experiment, both tracking and delivery systems are used. The tracking system provides users with the locations of particular colored targets whereas the delivery robot requires a location to deliver to. By linking these semantically similar services, a dynamic location delivery function is created. In this scenario, the delivery target is a person wandering around the environment.

The robots are configured the same as Experiment 1, except one of the tracking robots is now mobile (T3 in Figure 3). Although this robot can traverse the whole environment while tracking, it is confined to the corridor where the stationary trackers are located.

Once the systems and the infrastructure are connected, a single user logs on and selects the target tracking service to determine whether the target is currently being tracked. Once they are, the user selects the linking option from the service menu, which prompts for the data-providing service and the command service. These are selected as the applicable target tracking and delivery service respectively.

The operation of the infrastructure at this point is the same as if two separate services were selected. Both connect to the user client, which checks if the data service is part of a link. Since it is in this case, the information is sent to the command service's system client as well as being displayed to the user. The command service's system client parses the received information for the goal location. This effectively automates the user's manual goal location entry. Consequently, the delivery robot is constantly provided with the updated target location. In the cases when the target is lost or there is no detected location information in the data sent to the delivery robot, it proceeds to the last provided location. This helps deal with temporary communication failures and target loss.

Once the delivery robot is within a meter of the goal location, it halts and sends a message to the user that it is at its commanded location. If the target starts moving again, the new location information acts as a new goal location and the delivery robot continues to follow the target. In this experiment, the tracking information is sometimes provided by the mobile tracking robot, which resulted in a two-robot convoy following the target. This application of a linked service could easily be extended to generate a leader-follower behavior for multiple robots if they need to be relocated *en masse*.

Discussion

It is difficult to design a multi-human, multi-robot interaction mechanism that caters for interaction at the large and small scale. In addressing this issue, we assume that systems are mostly autonomous and interactions with them are loosely coupled. Our two-phased interaction infrastructure allows many systems to expose their services to many users, and upon selection, allows for more direct interaction between a single system and a user.

Our approach is focused more towards high-level interaction between users and systems rather than the tight coupling required for low-level interactions. This is partly due to the user interface requirements for more

complicated tightly-coupled interaction. In our approach, the interface is instantiated upon the user selecting the service and is based on the type of interaction. The more complex the interaction, the more complex the interface would have to be which can be difficult to generate *in situ*. Possible approaches to enhance the current interface are:

1. let the user configure their own interface from a set of generic components upon service selection (Dorohonceanu, Sletterink, and Marsic 2000);
2. let the infrastructure design the interface based on the semantics of the interaction (Ponnekanti et al. 2001);
3. let the system transfer a user interface program when the user selects a service from it. This is similar to the methods adopted by JINI to manage the user interface issue;
4. use the Internet as the interface medium since it allows many modes of interaction.

In any case, the interface may not provide the full modality and functionality that a system requires for tightly-coupled interaction. This is a reasonable compromise since it is postulated that a more general approach to the interface allows for a greater diversity of systems and services to be managed. Allowing this diversity is an important and necessary requirement for large-scale robotic projects, where it cannot be assumed that all robots are homogeneous and have the same interaction type. Nor can it reasonably be assumed that every system involved in such large-scale operations will be known in advance and have their desired interfaces available.

Another step to enhance the infrastructure is to distribute the server to provide greater redundancy. This will allow more robust operation in the presence of either communication or server failures. Currently, if the server fails, the infrastructure is unusable for new interactions, however any existing interaction between systems and users remains unaffected.

Acknowledgements

This work was supported by the DARPA MARS Program grant nos. DABT63-99-1-0015 and 5-39509-A, and DARPA SDR Program grant no. 4400057784.

Bibliography

Ali, K., and Arkin, R. 1994. Integration of Reactive and Telerobotic Control in Multi-agent Robotics Systems. In *Proceedings, From Animals to Animats 3: Simulation of Adaptive Behavior 1994*, 473-478, MIT Press.

Arkin, R., Collins, T., and Endo, Y. 1999. Tactical Mobile Robot Mission Specification and Execution. In *Proceedings, SPIE-99, Mobile Robots XIV*, 150-163, Boston, MA, September.

Dorohonceanu, B., Sletterink, B., and Marsic, I. 2000. A Novel User Interface for Group Collaboration. In

Proceedings of the 33rd Hawaii International Conference on System Sciences, Maui, Jan. 4-7.

Elhajj, I., Tan, J., Xi, N., Fung, W., Liu, Y., Kaga, T., Hasegawa, Y., and Fukuda, T. 2000. Multi-site Internet-based Cooperative Control of Robotic Operations. In *Proceedings of the 2000 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 826-831, Takamatsu, Japan.

Fox, A., Johanson, B., Hanrahan, P., and Winograd, T. 2000. Integrating Information Appliances into an Interactive Workspace. *IEEE Computer Graphics and Applications*, 54-65, May/June.

Gerkey, B., Vaughan, R., Støy, K., Howard, A., Sukhatme, G., and Matarić, M. 2001. Most Valuable Player: A Robot Device Server for Distributed Control. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2001)*, 1226-1231, Wailea, Hawaii, Oct. 29-Nov. 3.

Goldsmith, S., Spires, S., and Phillips, L. 1998. Object Frameworks for Agent System Development. Technical Report, SAND98-0742A, Sandia National Laboratories.

Jones, H., and Snyder, M. 2001. Supervisory Control of Multiple Robots Based on a Real-time Strategy Game Interaction Paradigm. In *Proceedings of the 2001 IEEE International Conference on Systems, Man, and Cybernetics*, 383-388, Tucson, AZ, October.

Jung, B., and Sukhatme, G. 2002. A Region-based Approach for Cooperative Multi-Target Tracking in a Structured Environment. forthcoming, *IEEE/RSJ International Conference on Intelligent Robots and Systems*.

Lee, J., and Hashimoto, H. 2000. Intelligent Space - Its concept and contents. In *Proceedings of the 2000 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2:1358-1363, Takamatsu, Japan.

Murphy, R., and Rogers, E. 1996. Cooperative Assistance for Remote Robot Supervision. *Presence*, Special Issue, vol. 5, no. 2, 224-240.

Phillips, L., Goldsmith, S., and Spires, S. 1998. CHI: A General Agent Communication Framework, Technical Report, SAND98-2825C, Sandia National Laboratories.

Ponnekanti, S., Lee, B., Fox, A., Hanrahan, P., and Winograd, T. 2001. ICrafter: A Service Framework for Ubiquitous Computing Environments. In *Proceedings of the Ubiquitous Computing Conference. Lecture Notes in Computer Science*, 2201:56-75.

Rybski, P., Burt, I., Dahlin, T., Gini, M., Hougen, D., Krantz, D., Nageotte, F., Papanikolopoulos, N., and Stoeter, S. 2001. System Architecture for Versatile Autonomous and Teleoperated Control of Multiple Miniature Robots. In *Proceedings of the 2001 IEEE International Conference on Robotics and Automation*, Seoul, Korea.

Surmann, H., and Theißinger, M. 1999. ROBODIS: A Dispatching System for Multiple Autonomous Service Robots. In *Proceedings of the FSR'99 Robotics Applications for the Next Millenium*, 168-173, Pittsburgh, PA.

Tews, A., Matarić, M., Sukhatme, G., and Gerkey, B. 2002. G'day Mate. Let me Introduce you to Everyone: An Infrastructure for Scalable Human-System Interaction. Technical Report, CRES-02-004, Center for Robotic and Embedded Systems, University of Southern California, August.

Tews, A., Sukhatme, G., and Matarić, M. 2002. An Infrastructure for Large-Scale Human-Robot Interaction. Technical Report, CRES-02-001, Center for Robotic and Embedded Systems, University of Southern California.