[Technical Reports (CIS)](...)   Department of Computer & Information Science

January 1993

Gigabit Telerobotics: Applying Advanced Information Infrastructure

Ruzena Bajcsy
University of Pennsylvania

David J. Farber
University of Pennsylvania

Richard P. Paul
University of Pennsylvania

Jonathan M. Smith
University of Pennsylvania, jms@cis.upenn.edu

Gigabit Telerobotics: Applying Advanced Information Infrastructure

Abstract

Advanced manufacturing concepts such as "Virtual Factories" use an information infrastructure to tie together changing groups of specialized facilities into agile manufacturing systems. A necessary element of such systems is the ability to *teleoperate* machines, for example telerobotic systems with full-capability sensory feedback loops. We have identified three network advances needed for splitting robotic *control* from robotic *function*: increased bandwidth, decreased error rates, and support for isochronous traffic. These features are available in the Gigabit networks under development at Penn and elsewhere.

A number of key research questions are posed by gigabit telerobotics. There are issues in network topology, robot control and distributed system software, packaging and transport of sensory data (including wide-area transport), and performance implications of architectural choices using measures such as cost, response time, and network utilization.

We propose to explore these questions *experimentally* in a joint research effort combining the Distributed Systems Laboratory (DSL) and the General Robotics and Sensory Perception (GRASP) Laboratory at the University of-Pennsylvania. The proposed experiments should provide important early results. A detailed research program is described.

Comments

Gigabit Telerobotics: Applying Advanced Information Infrastructure

Ruzena Bajcsy
David J. Farber
Richard P. Paul
Jonathan M. Smith
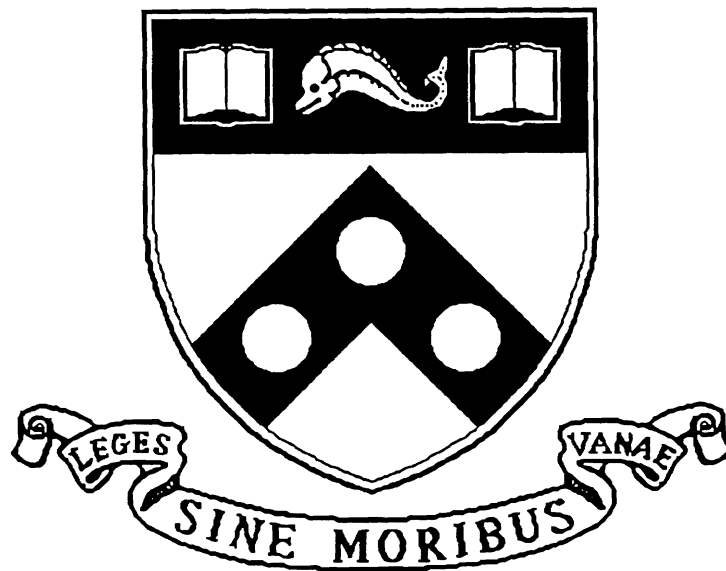
University of Pennsylvania
School of Engineering and Applied Science
Computer and Information Science Department

Philadelphia, PA 19104-6389

January 1993

Gigabit Telerobotics: Applying Advanced Information Infrastructure

Ruzena Bajcsy, David J. Farber, Richard P. Paul, Jonathan M. Smith

{jms,lou,farber,bajcsy}@cis.upenn.edu

Computer and Information Science, University of Pennsylvania
200 South 33rd St., Philadelphia, PA 19104-6389

ABSTRACT

Advanced manufacturing concepts such as "Virtual Factories" use an information infrastructure to tie together changing groups of specialized facilities into agile manufacturing systems. A necessary element of such systems is the ability to *teleoperate* machines, for example telerobotic systems with full-capability sensory feedback loops. We have identified three network advances needed for splitting robotic *control* from robotic *function*: increased bandwidth, decreased error rates, and support for isochronous traffic. These features are available in the Gigabit networks under development at Penn and elsewhere.

A number of key research questions are posed by gigabit telerobotics. There are issues in network topology, robot control and distributed system software, packaging and transport of sensory data (including wide-area transport), and performance implications of architectural choices using measures such as cost, response time, and network utilization.

We propose to explore these questions *experimentally* in a joint research effort combining the Distributed Systems Laboratory (DSL) and the General Robotics and Sensory Perception (GRASP) Laboratory at the University of Pennsylvania. The proposed experiments should provide important early results. A detailed research program is described.

1. Introduction

One of the major difficulties facing the United States is the inability to rapidly design, develop, and deploy complex systems, such as software and manufactured systems. Although our competitors on the world scene have made substantial progress in the ability to rapidly produce variants of established products, all nations seem to have difficulty in the creation and deployment of new systems. With the proposed project, the University of Pennsylvania's Computer and Information Science Department stands ready to develop technology that will help address these problems. This project represents an integration of disciplines that could, in the future, offer the same magnitude of payoff that the integration of computers and communications has. The core problem has been, and continues to be, one of *distributed resource management*.

A recent National Research Council study [1] on the United States Air Force's software projects found that a principle cause of the failure of the systems to achieve on-time and on-budget delivery was the inability to produce software systems in which pieces of the system were designed and built by different, geographically-distributed organizations. Similar problems arise in the integration of aircraft control systems into airframes and components of airframes with other components.

We propose applying the University of Pennsylvania's leading-edge work in robotics and manufacturing technology, together with our expertise in very high speed networking (as illustrated in the AURORA [5] project) towards the solution of these problems.

We are too realistic to promise great solutions immediately. We believe, however, that a reasonable start can be made and that short term benefits that will help our industrial base can be realized. Our proposal is described in the sections of this white paper. The unifying theme is that of a "Virtual Distributed Factory" (VDF). The VDF is geographically-distributed and multi-organizational [2]. It is composed of design centers [3], fabrication facilities,

and test stands, in addition to supercomputers for modeling, all integrated via a very high speed communication network. Integration of control and related actuation equipment (by humans or machines) requires *teleoperation*.

1.1. Teleoperation

Distributed digital teleoperation provides for the remote visual, audio, and kinesthetic access of an operator to a physical system. While many believe that visual access is all that is necessary, reflection (and experience!) reveals the difficulty of interpreting graphics displays and the impossibility of visually ascertaining how parts are connected together when partially mated. Kinesthetic feedback simplifies the manipulation of images and of parts. Combining audio information as well provides for an even stronger sense of reality. The use of more than one form of task feedback provides us with the capability to fuse sensory information, providing more complete knowledge of the state of a task and its performance. Kinesthetic feedback is a basic sense, and provides a very natural interface to systems. This interface is one reason that teleoperation is so direct and simple to use. In distributed digital teleoperation the operator interacts with a remote site via graphics displays, master manipulators, audio feedback, etc. There is nothing, however, in the system to tell the operator that she or he is interacting with a real remote system. We could replace the real remote world with a simulation. If we did this we would have created a virtual reality. While there is much research interest in virtual realities there is none, or very little, in distributed virtual reality. If a distributed digital teleoperation system existed it could be connected to a remote virtual reality allowing an operator to interact with products that have not yet been created.

In the area of modern manufacturing the rapidity of development, and correctness, of a new design is of utmost importance [3]. This is illustrated by CAD/CAM systems which are the forerunners of virtual reality systems. In the area of architectural design, virtual realities are used to allow users and architects to evaluate designs long before they are committed to paper. In manufacturing, we expect similar systems to become available; systems in which it is possible to "kick the tires" long before a product goes into detail design and prototype production. Such evaluation systems are of vital importance, as there is no longer time or money to go through the lengthy prototype and testing phases of the past.

This exploration and evaluation process can be expanded such that:

1. Multiple parts suppliers can test the assembly of the whole system electronically and make quick adjustments before the manufacturing process begins;

2. Physical constraints are made less relevant. For example, designers and production facilities may be in different physical locations, e.g., production is located in remote areas because real estate is cheaper (or zoning is less restrictive), while designers are typically concentrated in urban areas.

3. CAD systems are enhanced with physical simulators, such as stiffness analysis, thermo analysis, and fluid mechanics analysis. These simulators can be viewed as idealized knowledge representations of the physical processes. As such, they can predict the performance of the parts and/or assemblies under varied physical stresses. However, they can also, if connected to sensors, can be updated and refined when the actual system is built. This is particularly important when the same part is built with new materials and/or new manufacturing processes.

1.2. Design and Distributed Virtual Reality

Today we need to be able to conceive and build new products, without the need for a testing phase. This can be done by making use of virtual reality systems to ensure the correct function and manufacturability of the product, before it is actually built.

While these design problems can be solved by local virtual reality systems, in which visual, kinesthetic, and audio feedback are provided, there is also a need to develop complicated products, such as automobiles, which will be designed distributively with designers located in different countries all over the world. In order to do this we need to be able to interact with remote simulations of partial designs by methods identical to the distributed digital teleoperation we have described above. With such systems in place we would be able to provide for concurrent distributed engineering. We envision a system in which a product is designed "top down" and major components are given to design houses in locations where the best expertise is available. In order to ensure a valid final system design, each component designer will design in conjunction with virtual realities provided by all other component designers. At the same time each component designer will provide a virtual reality of their design. This virtual reality will become increasing detailed and accurate as a design progresses. At the same time that the functional

design is taking place, manufacturing engineers will be able to interact with the same virtual realities of the system to investigate its manufacturability. These manufacturing engineers will also be able to interact with the designers to modify a design to ensure manufacturability. It will be possible to concurrently design manufacturing equipment based on the same virtual reality information and to verify their operation. Finally, maintenance engineers will also be able to interact to ensure that a design is maintainable. In order to provide for the interaction between designers, located at many places in the world, distributed digital teleoperation systems will be necessary. A designer will need to have the ease of access to remote systems in which they can handle parts, operate subsystems, assemble sub-assemblies, verify access, etc. It is only with a combination of kinesthetic and visual feedback that this natural interface can be provided. Communications networks are just now becoming capable of supporting the appropriate feedback mechanisms.

1.3. Organization

The rest of this paper is organized as follows. The next section introduces teleoperation and sets the stage for a discussion of networking and new network capabilities in Section 3. Section 3 also suggests how the information infrastructure based on these network capabilities can be applied to Distributed Manufacturing. Section 4 lists the key networking research questions raised by this application. Section 5 gives our research plan, and Section 6 lists some expected results.

2. Telerobotics/Distributed Digital Teleoperation

Teleoperation is the performance of work at a distance. Initially conceived to handle radioactive materials, teleoperation allows an operator to exert forces on or impart motion to a slave manipulator and to experience kinesthetic feedback - that is, experience the forces on and motion of the slave manipulator. An operator is also provided with visual, and possibly audio feedback from the remote site. With the introduction of teleoperation technology, operators were able to perform sophisticated material handling and servicing tasks with considerable dexterity thanks to the natural interaction of the operator with the system. Even with limited visual and kinesthetic feedback, operators are able to accomplish tasks by using ingenuity and reasoning to compensate for the lack of sensory feedback and to overcome unexpected events.

2.1. Delay and Network Requirements

If we try to perform teleoperation at considerable distances, as, for example, in outer space, communications delays become significant. If the delay between an operator's actions and the feedback exceeds one third of a second the system becomes unstable and it is impossible to perform tasks using teleoperation. For electronic terrestrial communications, however, delays are within one third of a second, so it might be possible to provide for teleoperation between any two points on the earth's surface.

Unfortunately, teleoperation requires a number of communications channels, each of which has stringent requirements. Visual information communications requires megabit bandwidths with frame rates in excess of ten frames per second. Normally, teleoperation makes use of two to three video channels. The kinesthetic communications channel is required in both directions for each manipulator, of which there are typically two.

Such channels require the transmission of some hundreds of bits at the kilohertz rate; there are strict timing requirements on these channels and irregular, or missing data will seriously degrade system performance. Along with these channels might be channels for audio or tactile information. With the advent of modern digital wideband communications networks it appears possible to provide these communications channels digitally and to thus enable long distance teleoperation.

2.2. Impact

Long distance teleoperation would enable us to interact physically with any system from anywhere in the world. Experts in the servicing of equipment could provide emergency service from a central office to equipment located anywhere in the world - equipment such as nuclear reactors, aircraft, and computers. Equipment, located at many different places, such as remote transmitters and monitoring equipment, could be operated from a central location by a single operator. As the transmission of sensory data becomes more sophisticated, engineers could examine new products, prototypes and component parts. Bomb disposal experts could be available anywhere in the world they were needed while remaining in a secure central location.

Such a system would eliminate the need to replicate service personnel and experts all over the world and, furthermore, would provide for a higher level of training - fewer experts would be using their skills on a day to day basis, rather than waiting for an event to occur. Long distance teleoperation would also eliminate time delays and the need for travel and accommodation in moving experts from one location in the world to another. Productivity of experts would be increased due to the elimination of wasted travel time.

Much of the technology for long distance teleoperation is already in place, such as teleoperation and digital network communications. Work in our laboratory [11, 12] already provides the basis for digital teleoperation and similar work has been performed by Bejczy at JPL [4]. This work elevates communications between master and slave from the signal level to the symbolic. Instead of a stream of joint positions and torques, Cartesian displacements and forces are transmitted along with state identifying commands. In order to provide for distributed digital teleoperation we will need to insert a digital wide-band network between master and slave stations. We will need to establish communications protocols to provide for guaranteed real time communications. We will also need to determine what modifications are needed in the form of communications, used in teleoperation, to best utilize the advantages of the communications medium.

Once we have established a basis for distributed digital teleoperation a number of productivity improving technologies may be applied to teleoperation from the research area of robotics. Robotics is based on a geometric model of a task in which it is possible to reason about task performance and to automate much of the performance of a task. With distributed digital teleoperation, employing symbolic communication instead of signals, we may easily add a world model of a task and then apply robotics techniques directly. Such features as sub-task repetition, automatic tool change, automatic part insertion, collision detection may all be added to a system. Of importance here is that the operator always remains in control of the system, making use of automation where helpful, but always providing decision making and supervision of task execution.

3. Networking

In the past half-decade a number of interesting developments in the area of high-speed networking have taken place. In particular, we see the following:

A.   Use of fiber optic technology

B.   Availability of fast packet switches

C.   Network architectures which support isochronous (paced) traffic, such as voice and video

One result of these technological changes is the ability to provide network configurations with both selectable Quality of Service (QoS) characteristics and the high bandwidth necessary for transporting and reproducing sensory data at remote sites. The implication is that very complex systems, incorporating sensory data from the real world, can be combined with models, simulators and control paradigms from the computational world. Such systems might be applied to many tasks in which integrating distributed information and specialized resources can speed the process of turning concepts into real-world artifacts. For example, a two-hand robotic system with a visual input system might be used to explore, visually and kinesthetically, a remote system or virtual reality. Such systems can lead to rapid gains in manufacturing capability, as the rapid exploration possible with computational "realities" can be used to examine a larger design space of prototypes than would otherwise be possible.

Changes in networking have led to a number of proposals for the design of distributed multimedia communications systems. The challenges of these systems are the varied QoS requirements and the high bandwidth necessary for video. In the past, these requirements have led to the design of small-scale systems with specialized interconnection schemes, using slow (<100Kbits/sec) digital networks such as ISDN for voice and data communications, while using cable television for analog video services. The difficulty with this approach is its unwieldy use of specialized networks in order to meet the application service requirements. An additional difficulty with such systems is the potential need to add additional specialized networks whenever a new service with specialized QoS requirements is added. A more attractive approach integrates all media traffic into a single data stream; the data stream elements are characterized by their individual QoS requirements which are specified in a descriptive header. Thus, the addition of additional data types is accomplished using programmable software elements, rather than by laying new cables. We call this an 'integrated media'' approach to multimedia systems [9, 10]. In order to challenge this model, a non-traditional application was desirable. The close relationship between the DSL and GRASP laboratories at Penn led our group to the conclusion that "telerobotics" was an ideal application for both networking and robotics research.
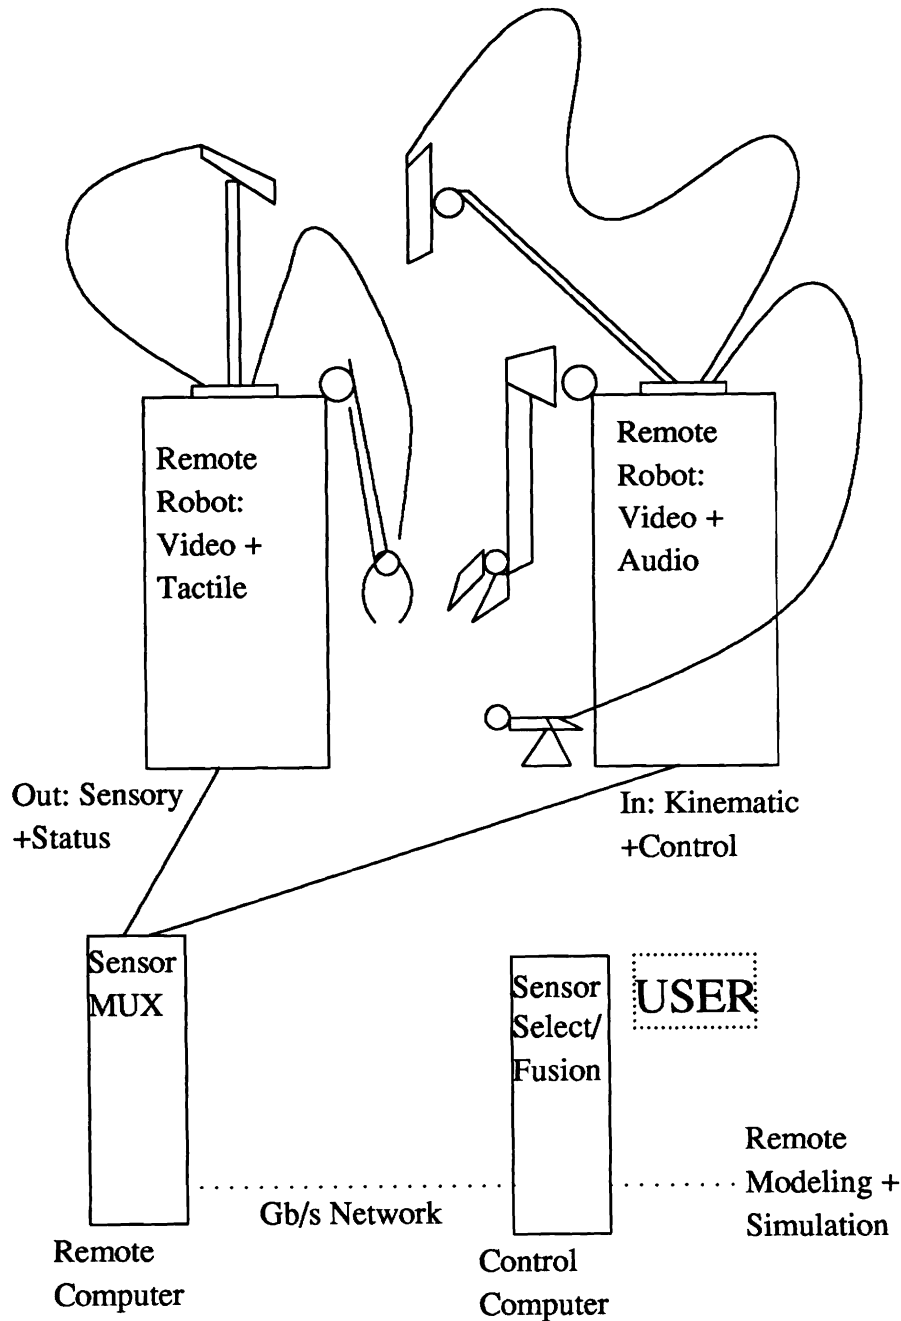
Figure 1: Gigabit Teleoperation

Another, and perhaps more important, consideration was the fact that any results have a broad impact due to their implications for advanced manufacturing systems.

3.1. Networking Implications of Telerobotics

The telerobotics application provides a high-bandwidth sensory stream, comprising (perhaps multiple) audio, video, kinematic and tactile data substreams. Component streams are gathered at a remote system, or perhaps from several distributed components which form a remote "aggregate" (e.g., a two-arm robotic system with shared sensory apparatus). Such an apparatus is illustrated in **Figure 1**. In the figure, a pair of robotic subsystems are combined in a

manufacturing task. Each of the two subsystems has a complement of sensory data-gathering equipment, such as video cameras, tactile sensors, and microphones. This delay-sensitive, high-bandwidth traffic is gathered by a remote computer, combined into a single stream, and carried to the computer labeled "Control Computer". As there is a finite limit on the bandwidth available for sensory data, there is great advantage to be derived from "intelligent" positioning and selection of different sensors. This "intelligence" is applied, depending on the tasks or subtasks, to optimizing the informativeness of sensory observation within the bandwidth limitations. This problem was addressed by Hager [8], but much work remains.

The stream of sensory data has varied QoS requirements and presents architectural challenges for the workstation software designer. It must be transported to a control point at which the data substreams are extracted, reproduced, analyzed, and used to direct commands back to the robotic system. While video can be noisy and tolerates dropped packets, the command streams are extremely loss- and delay-sensitive. The analysis of data at the "Control Computer" may involve advanced techniques for sensor fusion, in which tactile and video data are combined in order to categorize an object using extra data to determine otherwise hidden features. Selection may be done based on the focus of attention. The commands may be sent by another piece of apparatus - e.g., a master arm to which a particular remote arm is "slaved", or to a mobile sensory acquisition system (e.g., a set of cameras or other sensors) - to control the position of, or reposition, the data acquisition system. In this case, kinematic data is gathered from the master and used to generate a control and command stream for the robotic slave. The same principle applies to the control of the mobile data acquisition system. Other possible command generation scenarios include data gloves, mice, etc. One particular difficulty with the system as illustrated is the synchronization of sensory data and command streams for the two pieces of robotic apparatus. Note that **Figure 1** illustrates another possibility which only advanced networks offer, namely the ability to bring remote modeling and simulation resources (e.g., at specialized supercomputing sites) to bear on the manufacturing task, as needed. The gigabit per second bandwidths available allow sensor data from the robotic apparatus to be fed into the simulators and modeling software in real-time, allowing them to become components of an "intelligent" aggregation of components.

The particular features of advanced information infrastructure permitting such aggregates is discussed in the next section.

3.2. Advanced Networks

Typical real-time systems (such as those used for robotic control) have largely avoided the use of computer networks for their implementation. We see a number of reasons which might account for this, but the most important seems to be that common medium-speed Local-Area Networks, such as the Ethernet, are based on shared-bus media. Thus the degree of sharing, which is not known *a priori* can give rise to unpredictable network delays, which are unacceptable. Many other non-shared media either are too costly for general-purpose use or lack the bandwidth needed for the delivery of sensory data.

New network architectures based on high-speed packet switches have been developed. While a key design feature of these networks is their high (Gbps) bandwidth, they offer, in addition, bounded delay in switching components and bounded packet interarrival times. This allows real-time systems to be built from real-time systems connected by network fabrics which provide time-constrained services.

The particular examples being pursued in the AURORA testbed [5] are IBM's Packet Transfer Mode (PTM) and the Asynchronous Transfer Mode (ATM), which is the implementation proposed by the telecommunications industry for Broadband ISDN. The Distributed Systems Laboratory at Penn has been a strong contributor to this testbed, which has prototyped and explored a wide range of technologies. It is expected that some combination of these technologies will be employed in the development of a nationwide telecommunications infrastructure [7].

4. Research Questions

There are several key questions which this driving application gives rise to. In particular,

A.   Can effective hardware/software architectures for teleoperation be developed, and what are the environmental characteristics (such as errors, bandwidth fluctuations, and delays) which force system trade-offs?

B.   Can the widely-varying Quality-of-Service required for this application be reasonably accommodated by using a single fabric with high bandwidths?

C.   What software model is required in order to:

- allow specification of QoS
- map QoS specification to CPU scheduling and bandwidth allocations
- coordinate complex subsystems with varying (but strict) timing requirements in a general way
- provide extensibility in the face of new devices and challenges

D.   How will robotics be connected to future public networks, and what effects will switches and lengthy communications paths have on the behavior of distributed manufacturing systems connected with this infrastructure?

E.   Where should functions such as *sensor fusion* and *sensor selection* be embedded in the architecture of such a system? With a variety of options, what are the tradeoffs for capability and performance?

F.   How can simulations and models be improved with a real-time flow of real-world data? Does rapid model correction allow better decisions? Conversely, how does the improvement in modeling translate into reductions in demand for networking capability?

G.   How can we represent robotic/manufacturing tasks and subtasks? How can these representations be used to drive allocation and selection of manipulators and mobile sensory acquisition systems? Can we use the translation of the representation to real-world objects and actions not only to accomplish the tasks and subtasks, but to optimize our configuration with respect to time and other cost measures?

H.   What *language* can we use to describe manufacturing tasks? At what level of detail must objects be specified for manufacture or reliable "virtual" fitting and modeling? Can we generalize from MOSIS and CAD tools?

5. Research Plan

We propose a three phase plan in which the successive phases are increasingly challenging in their experimental requirements. Briefly, phase 1 will connect a pair of workstations co-located with the robotics apparatus (a master/slave pair of arms already present in our robotics laboratory) using electrical connections to realize a high-speed point-to-point networking fabric with low delay. During this phase the robotics devices (both manipulators as well as the data acquisition systems) will be interworked with the graphical simulator. We will embed these ideas in the framework of virtual worlds simulators. As shown in **Figure 2**, the graphical simulator attempts to model the actions of the real arm. The graphical simulator will represent not only the geometry of the robot's world but also some simulation of the dynamics of the physical world as the robotic agents interact with it. The physical interaction will model stiffness and inertia of objects, and later, heat propagation in the environment.

Phase 2 will replace the short distance electrical connection with a fiber optic link between the buildings containing the DSL and GRASP laboratories (there are about two city blocks separating the buildings which house our two laboratories). This connection will enable us to test the remote control of the robotic agents. Initially, this will involve the performance of a simple task, namely, grasping one object, first in an uncluttered environment, and then in an increasingly cluttered environment. The grasping will be aided initially by visual information with the camera system manually placed in order that it can fully observe (without occlusion) the object to be grasped. If time permits, we shall add acoustic information by first tapping the object before grasping. This acoustic information can give some sense of what the object's material composition and mass are before it is grasped. Such information can influence the grasping and lifting strategy. The next experiment will be to assemble or disassemble an object composed of three parts.

Phase 3 will connect the DSL/GRASP gigabit per second subnetwork into the AURORA gigabit testbed, introducing several new influences on the experiment. In particular, there are: packet delays due to a 300 mile round-trip latency; the presence of packet switches; and the potential for scale-induced phenomena such as congestion. At some time during either Phase 2 or Phase 3, we expect to introduce transport protocols optimized for this application. In this last phase, we shall increase the complexity of the manipulatory tasks as well as increase the complexity of the environment. A "complex" task, for example, would be the assembly or disassembly of an object composed of more than 8 or 10 components. The components are put together by a complex sequence of pulling, pushing, screwing, turning and sliding. The complexity of the environment means that the parts are varied in both shape and material.
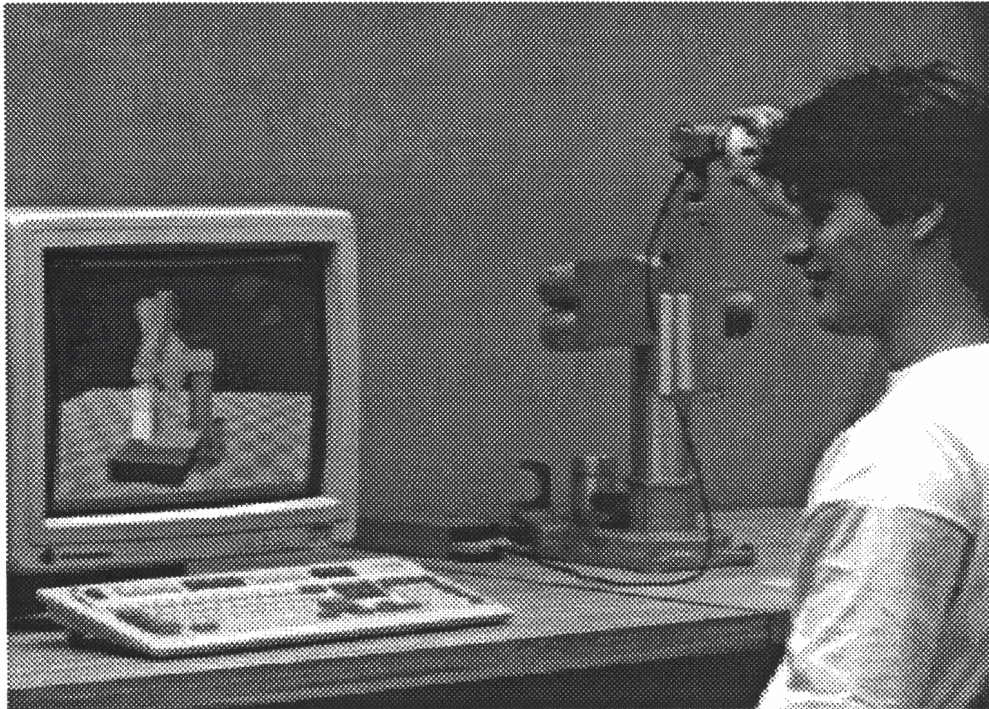
Figure 2: Robotic manipulator and model/simulation

At the end of these three phases (described in detail in the following sections), we would understand, and hopefully be able to demonstrate a realization of, remote performance of a complex task. This is the basis of distributed manufacturing, where control and realization, each using specialized facilities, will be separated. The complexity of the environment we describe encompasses environments with computerized controls, advanced programmable milling machines, and multipart assembly processes.

5.1. Phase 1

Phase 1 will connect two IBM RISC System/6000 workstations. The workstations will be connected using a pair of ORBIT [5] Micro Channel connected networking cards provided by the IBM T.J. Watson Research Center. Since the initial prototype will be used to answer feasibility questions and develop functional architectures, the IP protocol family will be used for data and command transport. While it is clear that IP is inappropriate for such transport in an actual full-scale network, the dedicated point-to-point nature of our network allows us to assume that IP will behave as if it was a real-time protocol. The main data volume is expected to be the video data gathered by the pair of cameras at the remote arm. This data will be captured by video capture hardware designed and implemented at Penn for the RISC System/6000 [14]. The captured video will be digitized and bundled into a sensory data packet along with tactile and other perceptual data for transport between the two ORBIT cards. The receiving system unbundles the data and presents it to devices suitable for reproduction, e.g., high-resolution color monitors. The gathered data are made available to a "teleoperator" who controls the small master arm. Kinematic data from the arm are captured and sent to the "remote" slave arm which is being observed. As mentioned above, if time permits we will add for Phase 1 an acoustic sensor which will receive information from tapping before grasping. This will provide further detail about the properties of the object, such as mass and material, which may improve both object identification and grasping performance. This work will serve as a driving application for operating system support needed for networked multimedia systems.

5.2. Phase 2

Phase 2 will replace the electrical (coaxial cable) connection between the ORBIT cards with a pair of fiber loops (ORBIT networks have a ring topology and are characterized, like other rings, by a bounded delay). In addition, an alternate topology using point-to-point ATM networks, implemented with a Penn-developed ATM Host Interface [13], will be deployed. During this phase we shall carry out the actual robotics trials, involving

assembly/disassembly experiments with full video, acoustic, and kinematic feedback. The issues of integration of sensing, task and subtask representation, control of the manipulators, and positioning of the sensors will be fully investigated covering a broad range of tasks, graduated from simple to complex tasks and environments.

5.3. Phase 3

Phase 3 will realize connectivity using a network connection which can be varied between plaNET/PTM networks and Sunshine/ATM. Each of these networks spans a 300 mile distance, as they connect the University of Pennsylvania, Bell Communications Research, the IBM T. J. Watson Research Center and the Massachusetts Institute of Technology. Each network includes packet switches and interconnects large numbers of powerful workstations. The characteristics of this network more closely represent the expected structure of future national public network infrastructure, and thus provide a fine framework for developing software which should have a long lifetime and great applicability.

6. Expected Results

We envision a number of important results from this research.

From the networking focus, we will begin to understand software/hardware architectures for combining the sensory data and the kinematic modeling in a system which is physically distributed. This includes characterizing (and developing frameworks for evaluating) the Quality of Service requirements imposed by a system with such rich sensory data, timeliness, and reliability constraints. Another natural outcome is an instantiation of these requirements into a protocol suite and distributed multimedia communications system which can use the features of advanced networking infrastructure effectively. A final networking result will be the understanding of the effects increasing latency and raising communications bandwidth will have on telerobotics *systems*.

A key result from the robotics focus will be the understanding of task complexity and environment complexity where autonomous systems fail but human augmentation allows successful performance. We know for example that visual segmentation of a general scene is an unsolved problem, since the desired segmentation and grouping is very task and context dependent. Here, we have a chance to determine the boundaries and limitations, both without a human, and with a remote human teleoperator.

Another key result will stem from from the determination of the types and volumes of sensory data a system must have for tasks involving assembly and disassembly. We expect that this information will be quantized with respect to the complexities of the task/operations and the environment.

Finally, we will explore a scenario which will give us an understanding of the effects of bridging distance via electronic means, applied to manufacturing. In a modern world with widely separated designers, suppliers and assembly plants, manufacturing is an ideal application for an advanced information infrastructure.

The importance of such systems research *cannot be overemphasized* - it forces you to get the details right. To quote R. Feynman, "*For a successful technology, reality must take precedence over public relations, for Nature cannot be fooled.*" [6] As the pace of competition increases, the spoils will go to those who gain the most leverage from their existing advantages to develop new advantages. We believe that an opportunity for tremendous leverage exists today in "Gigabit Telerobotics."

7. Acknowledgments

Craig Sayers created Figure 2, Joseph Farber provided constructive editorial comments, and Klara Nahrstedt helped to understand QoS.

8. References

[1]   *Adapting Software Development Policies to Modern Technology,* National Academy Press, Washington, DC (1989). Air Force Studies Board, Commission on Engineering and Technical Systems, National Research Council

[2]   "Networks Help Small Companies Think and Act Big - Members Join Forces to Compete With Larger Firms in U.S. and Abroad," *The Wall Street Journal*, p. B2 (November 12, 1992).

[3]   "Moving Past the Assembly Line: 'Agile' Manufacturing Systems May Bring a U.S. Revival," *Business*

Week, pp. 177-180, McGraw-Hill, Special Bonus Issue - Reinventing America: Meeting the Challenges of a Global Economy (October 23, 1992).

[4] A. K. Bejczy, W. S. Kim, and S. Venema, *The Phantom Robot: Predictive Displays for Teleoperation with Time Delay*, May 1990.

[5] D. D. Clark, B. S. Davie, D. J. Farber, I. S. Gopal, B. K. Kadaba, W. D. Sincoskie, J. M. Smith, and D. L. Tennenhouse, "An Overview of the AURORA Gigabit Testbed," in *Proceedings, INFOCOM 1992*, Florence, ITALY (1992).

[6] Richard P. Feynman, *What do YOU Care What Other People Think?*, Norton (1988), p. 237.

[7] Al Gore, "Infrastructure for the Global Village," *Scientific American* **265**(3), pp. 150-153 (September 1991).

[8] Gregory D. Hager, *Active Reduction of Uncertainty in Multi-Sensor Systems,* Computer and Information Science, University of Pennsylvania (1988). Ph.D. Thesis

[9] Klara Nahrstedt and Jonathan M. Smith, "An Integrated Multimedia Architecture for High-Speed Networks," in *Proceedings, Multimedia '92 Conference*, Monterey, CA (April 1992).

[10] Klara Nahrstedt and Jonathan M. Smith, "An Integrated Multimedia Architecture for High-Speed Networks (Brief Extended Abstract)," in *ACM Computer Communications Review* (July 1992).

[11] Richard P. Paul, Janez Funda , Thomas Lindsay , Simeon Thierry, and Masahiko Hashimoto, "Model Based Teleoperation to Eliminate Feedback Delay," MS-CIS-91-02, Univ. of Penn., Philadelphia, PA (January, 1991). NSF Grant BCS89-01352, First Report

[12] Richard P. Paul, Janez Funda, Thomas Lindsay, Craig Sayers , and Masahiko Hashimoto, "Model Based Teleoperation to Eliminate Feedback Delay," MS-CIS-92-75, Univ. of Penn. (October, 1992). NSF Grant BCS89-01352, Second Report

[13] C. Brendan S. Traw and Jonathan M. Smith, "Hardware/Software Organization of a High-Performance ATM Host Interface," *IEEE Journal on Selected Areas in Communications (Special Issue on High Speed Computer/Network Interfaces)* **11**(2) (February 1993).

[14] Sanjay K. Udani, *Architectural Considerations in the Design of Video Capture Hardware,* University of Pennsylvania, School of Engineering and Applied Sciences (April, 1992). M.S.E. Thesis (EE)