N93-32103

Stanford Aerospace Robotics Laboratory Research Overview

W. L. Ballhaus [*] L. J. Alder [†] V. W. Chen [‡] W. C. Dickson [§]
M. A. Ullman [¶]
Stanford University Aerospace Robotics Laboratory
Stanford, California 94305

Abstract

Over the last ten years, the Stanford Aerospace Robotics Laboratory (ARL) has developed a hardware facility in which a number of space robotics issues have been, and continue to be addressed. This paper reviews two of the current ARL research areas: navigation and control of free flying space robots, and modelling and control of extremely flexible space structures.

The ARL has designed and built several semi-autonomous free-flying robots that perform numerous tasks in a zero-gravity, drag-free, two-dimensional environment. It is envisioned that future generations of these robots will be part of a human-robot team, in which the robots will operate under the task-level commands of astronauts. To make this possible, the ARL has developed a graphical user interface (GUI) with an intuitive object-level motion-direction capability. Using this interface, the ARL has demonstrated autonomous navigation, intercept and capture of moving and spinning objects, object transport, multiple-robot cooperative manipulation, and simple assemblies from both free-flying and fixed bases.

The ARL has also built a number of experimental test beds on which the modelling and control of flexible manipulators has been studied. Early ARL experiments in this arena demonstrated for the first time the capability to control the end-point position of both single-link and multi-link flexible manipulators using end-point sensing. Building on these accomplishments, the ARL has been able to control payloads with unknown dynamics at the end of a flexible manipulator, and to achieve high-performance control of a multi-link flexible manipulator.

1 Experiments in Autonomous Navigation and Control of Multi-Manipulator, Free-Flying Space Robots

1.1 Introduction

Although space presents an exciting new frontier for science and manufacturing, it has proven to be a costly and dangerous place for humans. It is therefore an ideal environment for sophisticated robots capable of performing, as part of a human-robot team, tasks that currently require the active participation of astronauts.

As our presence in space expands, it will be increasingly important for robots to be capable of handling a variety of tasks ranging from routine inspection and maintenance to unforeseen servicing and repair work. Under the task-level guidance of astronauts, such tasks could be carried out by free-flying space robots equipped with sets of dexterous manipulators. These robots will need to be able to navigate to remote job sites, rendezvous with free-flying objects, perform servicing or assembly operations, and return to their base of operations.

In order to advance the underlying theory and technology necessary for the aforementioned robotic capabilities to be realized, the ARL has identified and addressed the problems associated with building and controlling autonomous

*Ph.D. Candidate, Department of Aeronautics and Astronautics.
†Ph.D. Candidate, Department of Aeronautics and Astronautics.
‡Ph.D. Candidate, Department of Electrical Engineering.
§Ph.D. Candidate, Department of Aeronautics and Astronautics.
¶Ph.D. Candidate, Department of Aeronautics and Astronautics.

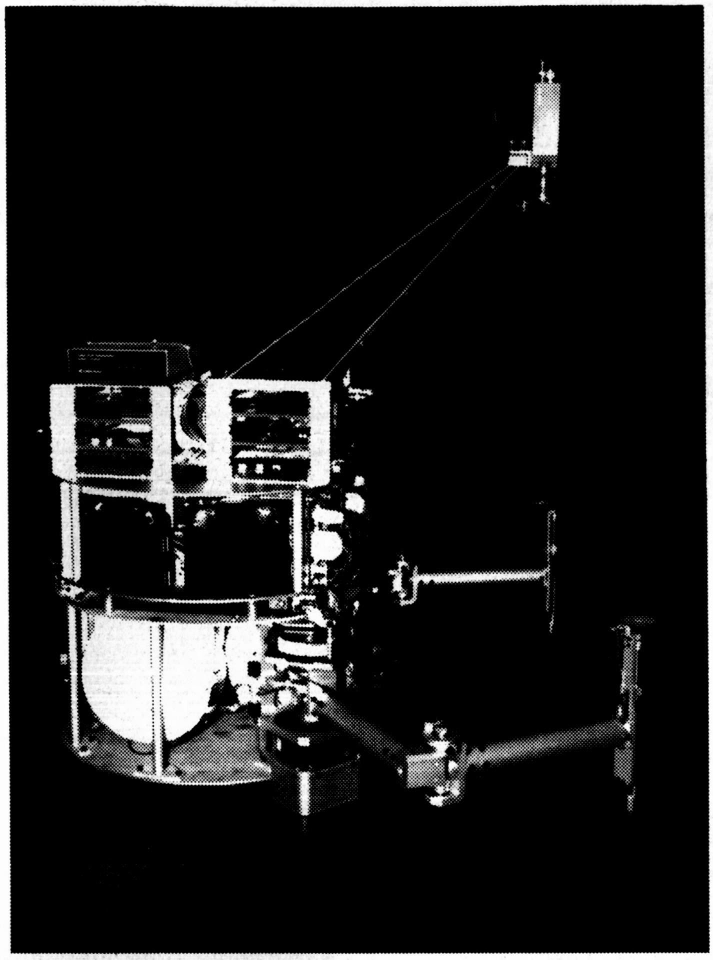

Figure 1: **Stanford Multi-Manipulator Free-Flying Space Robot**
This is a fully self-contained 2-D model of a free-flying space robot complete with on board gas, thrusters, electrical power, computers, camera, and manipulators. It exhibits nearly frictionless motion as it floats above a granite surface plate on a 0.005in thick cushion of air.

free-flying space robots. The objective of this research has been to demonstrate the ability to carry out complex tasks including acquisition, manipulation, and assembly of free-floating objects based on task-level commands. The approach has been to extend earlier ARL work in cooperative manipulation involving the use of fixed-base manipulators[1] to accommodate an actively mobile base thereby removing the workspace limitations inherent in fixed-base implementation.

1.2 Experimental Hardware

To test newly developed design methodologies and control strategies, the ARL has developed an experimental two-armed satellite robot shown in Figure 1. The robot uses an air cushion support system to achieve—in two dimensions—the drag-free, zero-g characteristics of space. The robot is a fully self-contained spacecraft possessing an on board gas supply for flotation and propulsion, rechargeable batteries for power, and on-board computers with sensing and driver electronics for navigation and control. Although the robot can function autonomously, its computers can also communicate with a network of workstations via a new wireless LAN. An on-board camera provides optical endpoint and target sensing while an overhead global vision system facilitates robot navigation and target tracking.[1] The robot "floats" on a $9' x 12'$ granite surface plate with a drag-to-weight ratio of about 10^{-4} and gravity induced accelerations below $10^{-5}g$—a very good approximation to the actual conditions of space. A more detailed description of the space robot is given in [2].

[1] The global vision system serves as a convenient laboratory surrogate for a tracking system such as GPS that could be used for this purpose in space.

1.3 Controller Architecture and Graphical User Interface

The controller architecture consists of a three-level hierarchy composed of a stateless remote graphical user interface, a high-level strategic controller, and a low-level dynamic controller based on an operational space computed torque formulation.

The graphical user interface (See figure 2) runs on a Sun Workstation and allows an operator to send high level commands to the robot by selecting icons and clicking on buttons.

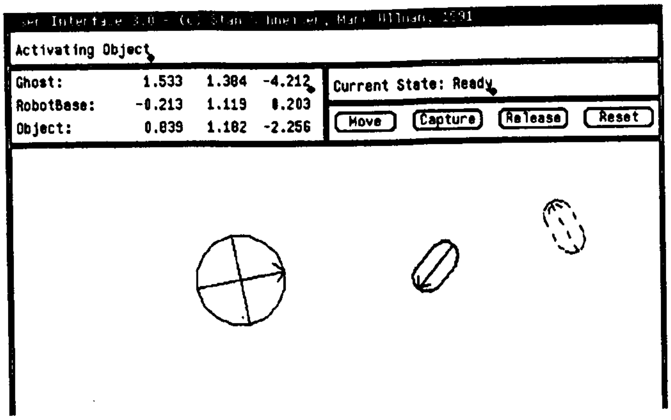

Figure 2: **Typical View of the User Interface**

The graphical user interface provides "point and click" operation of the robot and allows the operator to control and monitor all operations remotely. Here a capture and move operation is underway.

The high-level strategic controller is based on a sophisticated finite state machine. It accepts commands from the remote user interface and reconfigures the low-level dynamic controller to carry out desired actions. A thorough discussion of the strategic controller is beyond the scope of this paper and can be found in [2].

1.4 Object Rendezvous and Capture

The execution of useful work in space requires the ability to simultaneously control both robot base and manipulator motions. In general, rendezvousing with and capturing a free-flying object requires controlling both manipulator and base body positions to follow *coordinated* intercept trajectories. Global navigation and control (or "gross motion" control) of a space robot therefore poses a set of interesting and unique challenges. These differ fundamentally from both the typical satellite positioning/attitude control problem and the case of a free-floating space robot with an uncontrolled base.

The robot can be commanded to capture and retrieve an object via the graphical user interface described in the section *Controller Architecture and Graphical User Interface.* Figure 3 shows the time history of one such rendezvous. The object floats on an air bearing supplied by a battery-powered aquarium air pump. The object can be sent across the granite table in a random direction with a rotation rate as high as 20 revolutions per minute. Since the object is of comparatively low mass ($\sim 1kg$), the robot follows a straight line path to intercept it. Upon grasping the object by inserting its "peg-in-the-hole"-style grippers, the robot brings the object smoothly to rest (in the robot's reference frame). It can then stow the object and transport it to some new location where it can release it.

1.5 Nonlinear Adaptive Control of a Free-Flying Space Robot

Adaptive control for robots is useful in several important, common situations: 1) When there is poor or no knowledge of the payloads, 2) When there are inaccurate models of the robot, 3) When there are changes in the environment. While a robust, nonadaptive, controller may provide the same protections as an adaptive controller in these situations, it typically does so with substantially reduced performance. The added complexity of an adaptive controller wins back that lost performance.

The most obvious situation in which to use adaptive control is for handling payloads that have unknown or poorly known physical properties—for example, when handling damaged satellites where the nature and extent of the damage

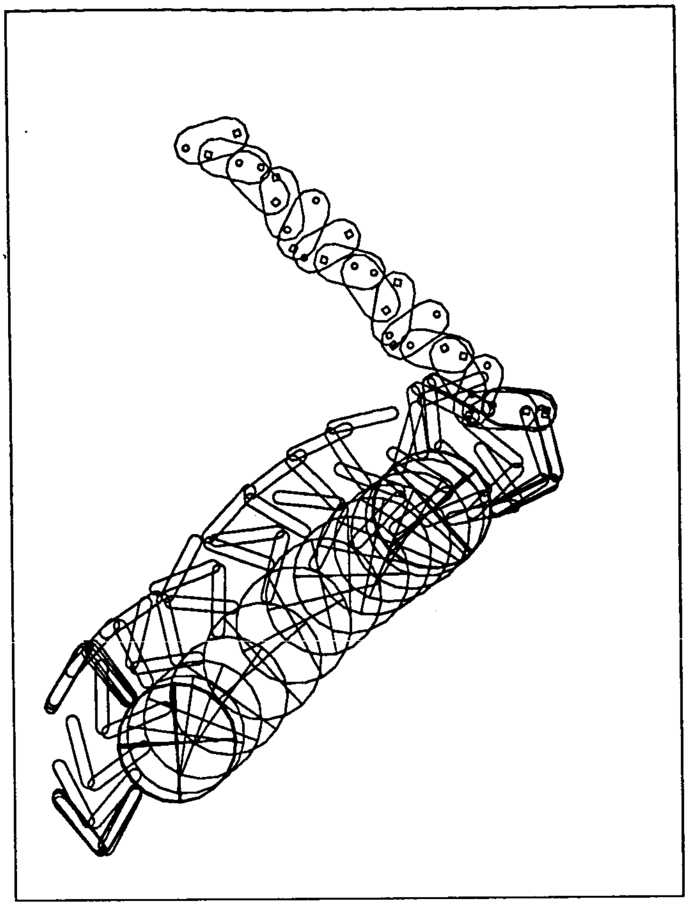

Figure 3: **Time history of Object Rendezvous and Capture**

This "time-lapse" plot of experimental data shows the motion of a spinning object and the path the robot executed in order to intercept and capture it. The frame rate is 0.5Hz and this figure shows about 30 seconds of elapsed time.

is unknown. More generally, this capability relieves astronauts of the duty to inform the robot of the detailed physical properties of each payload the robot is to handle. While a comprehensive parts database can relieve much of this responsibility, adaptive control provides protection in cases when the database is not completely accurate or is lacking.

Adaptive control also eases the basic controller design process. There are typically many aspects of the robot itself that are either poorly modelled or not modelled. It is difficult to develop accurate models for robots. While lengths and masses can be measured relatively accurately, effective center of mass locations and moments of inertia of individual links after being incorporated into the robotic system are estimates at best. Other poorly characterized effects, such as joint friction, and forces caused by wiring harnesses, contribute to an inaccurate model. By providing appropriate adjustable parameters to the controller, adaptive control can adapt to these uncertainties to render their effects unimportant.

This research has generated an adaptive control framework that is very general and easily extensible to even larger, more complex systems than the free-flying robot with cooperating manipulators for which it was developed [3]. The contributions made by this research include:

- Development of a general adaptive control framework—the adaptive *task*-space framework—that is capable of providing full adaptation to a free-flying space robot with two cooperating manipulators in all modes of operation.

- Extension and generalization of a joint-space, nonlinear, adaptive control algorithm, based on inverse-dynamics, to control in the *task* space, which represents a broader class of control inputs, including, but not limited to, cooperative *object* control, as well as endpoint control and joint control.

- Formulation of the *system concatenation* approach for efficient, incremental generation of system models for multiple, interacting systems. *System concatenation* takes full advantage of models already developed for each manipulator or robot subsystem to minimize the effort in deriving the total system models used for adaptation.
- Full integration of the new adaptive control algorithm into a hierarchical control architecture that includes a graphical user interface and a finite state table programming environment.
- Experimental verification of the new adaptive controller, in the hierarchical control environment, on a free-flying space robot model.
- Development of the "Point-Grabber II" vision system that, together with software drivers developed in ARL, is capable of tracking bright spots at 60 Hz with better than 1/20 pixel resolution.

1.6 Cooperative Object Manipulation by Free-Flying Robot Teams

Free-flying space robots could perform many of the dangerous and expensive extra-vehicular activity (EVA) tasks presently requiring humans, such as the assembly and maintenance of structures and off-board platforms. Certain missions may require two or more robots to work as a team to cooperatively manipulate large objects such as satellites or structural elements. Situations requiring this capability include the initial installation or assembly of objects, or the retrieval of objects for repair or replacement. Manipulating flexible or multi-body objects is another task which may require a team of robots—a single robot may not be able to adequately control the object's internal degrees of freedom.

With this motivation in mind, the overall objective of this research was to identify and address the central dynamics and control issues relating to object manipulation by free-flying robot teams [4]. In achieving this goal, this research made the following contributions to automatic control and robotics:

- The concept of a team manager was developed as a mechanism for directing the activities of multiple independent robots into a cooperative team effort.
- Task-level direction was extended for use with free-flying robot teams for the first time. Once the user specifies a desired object position, the robot team determines and executes the control commands required for proper manipulation.
- A *Hybrid-Dynamics* formulation of equations of motion was developed that determines the mixed set of system accelerations and controls that is consistent with a specified complementary mixture of accelerations and controls. This formulation provided the basis for a robot control algorithm that combines discrete-valued base thrusters with proportional arm motors to produce precise manipulator endpoint accelerations.
- A dynamic modelling method has been developed for producing a system model directly from subsystem descriptions. This method is based on the solid foundation of constrained-system theory and takes advantage of simple kinematic relationships to merge subsystem descriptions into a full system model with no need for subsystem decomposition.
- Laboratory experiments have successfully demonstrated a team of free-flying robots capturing and manipulating a large, freely moving object. In these experiments, a human user indicates a desired object location and orientation through a graphical user interface. The robots then capture and so position the object with no additional input required from the user.

2 Control of Flexible Space Manipulators

2.1 Introduction

Many current and future space missions either require now or will require the assistance of large scale manipulator systems. For example, the Shuttle Remote Manipulator System (SRMS) has aided in the deployment, maneuvering, and retrieval of satellites and orbiter payloads. Future space robotic tasks will include spacecraft inspection and maintenance, transportation of payloads about space structures, and docking maneuvers.

To maximize the workspace of space robotic manipulators, the links of these robots tend to be very long in length. Furthermore, due to the cost of boosting mass into orbit, lightweight manipulator designs are most desirable. The resulting large, lightweight space manipulators contain low-frequency inherent structural flexibility which increases the difficulty of achieving high-performance control. As a result, our ability to model and control these flexible manipulators will determine how efficiently they can be used to perform various tasks.

Large flexible space robots like the SRMS are currently controlled by astronauts using joy sticks to control the position of each manipulator joint. To avoid exciting the low-frequency modes of the SRMS, the astronauts move the SRMS very slowly so that the end-point speed remains typically below 0.06 m/sec and 0.6 m/sec in its fully loaded and unloaded configurations respectively.

Previous ARL research has demonstrated that the performance of large lightweight manipulators can be significantly increased by improving the system controller. Through the use of direct end-point sensing, the end-point position can be used in a feedback control strategy to achieve accurate control of the manipulator end-point [5] [6].

Current ARL research in the control of flexible space manipulators extends this earlier work to control payloads with unknown dynamics at the end of a flexible manipulator, and to achieve high-performance control of a multi-link flexible manipulator.

2.2 Control of a Flexible Robotic Manipulator with Unkown Payload Dynamics

Space-based robots such as the SRMS and the proposed Space Station Remote Manipulator System (SSRMS) have been and will be essential elements of future space exploration. Further, the payloads manipulated by these large, flexible robots may themselves include unknown internal dynamics. Examples include sloshing fuel and/or flexible appendages (e.g. vibrating solar panels on a small satellite). If the payload dynamics are not accounted for in the control design, degraded performance or instability are possible.

The existing body of control that has been developed for flexible space robotic manipulators is insufficient to achieve high-performance control of such complex configurations. Collocated controllers are insensitive to payload dynamics but are low performance [7]. High-performance control has been achieved by using non-collocated end-point control [7], but the configurations studied all involved payloads that were well modelled by a tip inertia matrix only. Further, it has been shown that these high performance non-collocated controllers are sensitive to the mass and inertia of the tip [8]. It has also been demonstrated that a controller that is tuned for a particular tip mass may in fact be unstable for a different tip mass [8]. The sensitivity to an unknown tip mass has been accounted for successfully using adaptive endpoint control [9].

The goal of this research is to develop control techniques that provide precise, high bandwidth end-point control of flexible manipulators and are also able to damp any internal oscillations of the payload. The internal dynamics of the payload will not be known a priori. Furthermore, it is assumed that it will not be practical to outfit the payload with sensors that measure the internal state of the payload. The sensory input for controlling the robot-payload system will be based on the robot system sensors only. [2]

The control approach that has been developed and demonstrated experimentally in this research is based on extensions to the self-tuning regulator solution of adaptive control (see [10] for details). The extensions yield a feasible, real-time control that potentially can be implemented in future space robotic systems. High performance control is merged with an innovative identification algorithm in a self-tuning regulator approach. The identification of the payload is done using recently developed subspace fitting techniques. These techniques allow real-time determination of the order of the payload dynamics. Sufficient excitation problems are addressed by performing the identification closed loop.

Figure 4 shows a schematic diagram of the hardware. The experiment represents a large space-based manipulator holding a payload that has internal vibrating dynamics. The robot arm is a flexible beam which moves in the horizontal plane. At one end of the beam, there is a motor, and on the other end there is the payload which is a pad that floats on an air bearing. The pad and air bearing prevent out-of-plane vibrations. Mounted on the pad is a pendulum that represents the dynamics of the payload. Floating the pad on a smooth granite table simulates the zero-g environment of space in one dimension.

The pendulum inside the payload can be held mechanically so that it will not vibrate, or it can be set free to oscillate. The length of the pendulum can be modified to simulate an unknown frequency of oscillation (i.e. one that might be encountered in a sloshing fuel tank). The damping ratio of the pendulum is on the order of 0.5%. When released form a 45^o initial condition, it takes about one minute to damp within $\pm 5^o$ of its steady-state value. Reference [11] has a more detailed description of the hardware including mass properties.

The effect of having unmodelled payload dynamics on the closed-loop performance of the system can be seen by examining the time response shown in Figure 5. Figure 5 shows the time response of an end-point-based LQG controller

[2] This includes vision sensing.

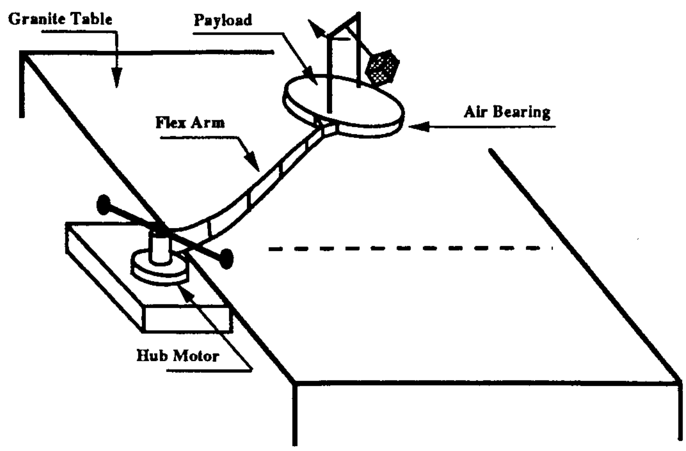

Figure 4: **Hardware Schematic**

This figure shows a schematic of the single link flexible manipulator with a dynamic payload.

that neglects the dynamics of the payload (i.e. the payload is considered to be a rigid body). This inadequate time response verifies that the payload dynamics cannot be neglected when using end-point control.

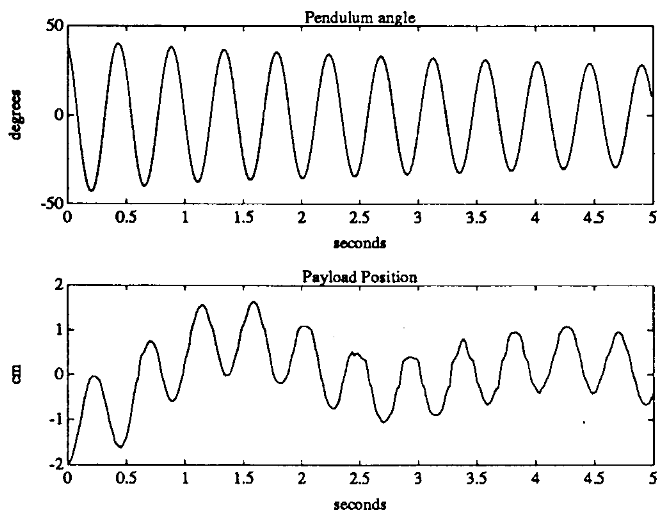

Figure 5: **Time Response**

This figure shows the time response of an end-point-based LQG controller that neglects the dynamics of the payload. This inadequate time response verifies that the payload dynamics cannot be neglected when using end-point control.

The experimental results of the adaptive controller are shown in Figure 6. The pendulum is mechanically held at a 45^o initial condition so that the payload is initially a rigid body. While the controller is regulating the payload position, the pendulum is released. For the first 2.2 seconds, the pendulum is being held at 45 degrees and the payload is regulated to zero. At 2.2 seconds, the pendulum is released resulting in motion of the payload. For the next two seconds, the pendulum is damping very slowly with a damping ratio of less the 0.5%. During this time, the identification algorithm detects the frequency of the pendulum, and at four seconds into the run, it swaps in a controller that accounts for the payload dynamics. Three seconds later, the pendulum is damped to within five degrees and the arm is within two millimeters of its desired position.

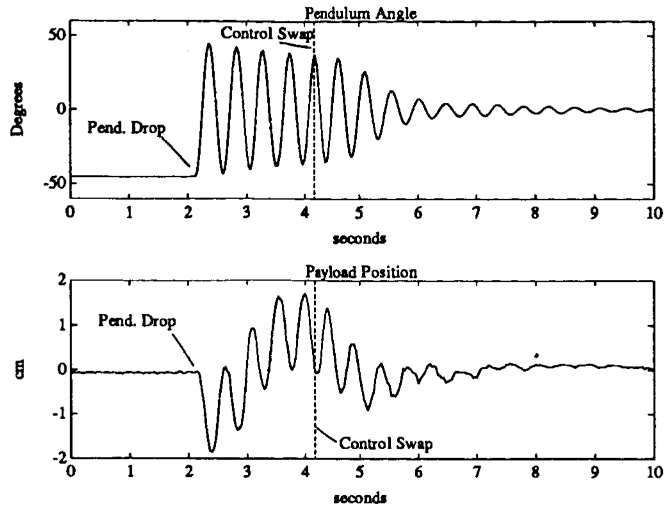

Figure 6: **Time Response**

This figure shows the time response of the adaptive controller. The identification algorithm detects the frequency of the pendulum, and then swaps in a controller that accounts for the payload dynamics.

2.3 High-Performance Control of a Multi-Link Flexible Manipulator

This research focuses on the problem of repositioning quickly and accurately a multi-link flexible space structure within a specified time (see [12] for details). It is assumed that only the initial and final states of the structure and the final time of the slew are given. A quadratic performance index is formulated which weights the value of the state at the terminal time, and the integral of the square of the control effort expended to that time. Minimizing this performance index with respect to the control effort results in a terminal controller with time-varying feedback gains.

For many applications, the motion of the end-point of a structure may be arbitrary, as long as the structure reaches a desired configuration within a specified time. For example, suppose a large flexible space robot was initially performing a task in one portion of its workspace, and then was commanded to perform a task at another location. In this case, assuming that the workspace is clear of obstacles, the trajectory that the endpoint follows in moving from its original working location to the commanded position is not of concern, as long as the manipulator reaches its desired destination, or repositions itself, quickly. Similarly, the transportation of rigid payloads from one location to another is an additional application where only the initial and final positions and the time duration of the slew are of importance.

There are various approaches that can be taken to reposition a flexible structure. For example, when a desired trajectory is known, it is possible to calculate an open-loop set of control inputs which would result in the desired end-point motion. Another method of repositioning a flexible structure or manipulator involves using full state feedback based on optimally determined, time-invariant (TI) feedback gains derived from a linear-quadratic cost function. This approach is fully described in [13] and [6]. Although these methods have been shown to perform quite well, they require a specified trajectory from which the control inputs are calculated. The terminal controller design discussed in this research automatically generates, based on the minimization of a linear-quadratic performance index, the optimal state history and the corresponding feedback gains necessary to enforce this motion. Terminal controllers have the added advantage of being computationally inexpensive to calculate (allowing for on-line controller design), and are independent of the initial and final states of the system. Furthermore, the time duration of the repositioning is the only slew parameter upon which the terminal controller design is dependent. As a result, the same set of controller gains can be used for point-to-point slews having equal specified final times.

The Stanford Multi-Link Flexible Manipulator, shown in Figure 7, consists of a two-link flexible manipulator which operates on air cushions in the horizontal plane of a 1.2 m by 2.4 m granite table. Each of the flexible links is 0.52 m in length, and consists of an aluminum beam (cross section 1 mm by 38.1 mm) with discrete masses evenly spaced along its length. The discrete masses, termed "Mass Intensifiers", increase the overall beam mass without changing its flexural rigidity. These Mass Intensifiers also lower the natural frequencies of vibration. The flexible links exhibit significant bending in the horizontal plane due to their narrow cross section and orientation. Air cushion supports at the elbow

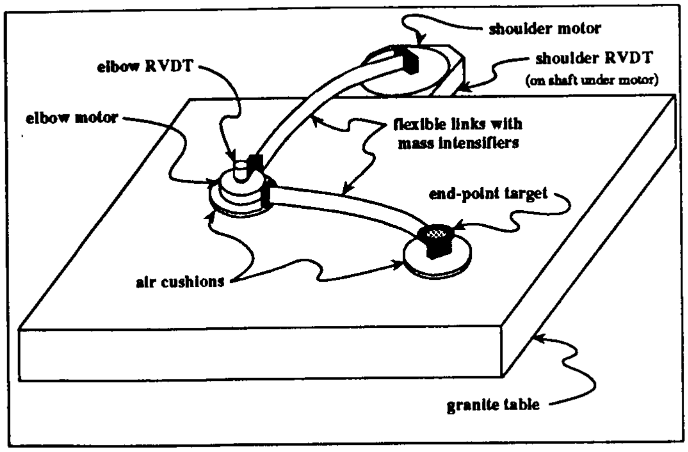

Figure 7: **Schematic of the Stanford Multi-Link Flexible Manipulator**

Shown in this figure is the Stanford Multi-Link Flexible Manipulator along with the various components of the experimental apparatus.

and end-point provide torsional stiffness. In addition to exaggerating the link flexibility, the structure was designed to be lightly damped. As a result, damping must be provided to the system through active control. The shoulder motor, mounted on the side of the granite table, can provide a peak torque of 5.43 N-m. The elbow motor, mounted on the elbow air-cushion pad, can provide a peak torque of 1.06 N-m. Both actuators are direct-drive, DC limited-angle torquers. Rotary variable differential transformers (RVDT's) are located at each of the motor shafts and provide joint angle measurements. A vision sensor, fully described in [1], is available for end-point measurements. It consists of a CCD television camera that tracks a special variable reflectivity target located at the manipulator end-point. The vision system has the capability to track multiple targets at a sample rate of 60 Hz with a resolution of approximately 1 mm over the roughly 1.5 m^2 workspace [6].

Figures 8 and 9 show a comparison between experimental and simulated responses for the terminal controller for a final slew time of $t_f = 2.5$ seconds. From Figure 8 it is evident that both the shoulder and elbow motor position histories

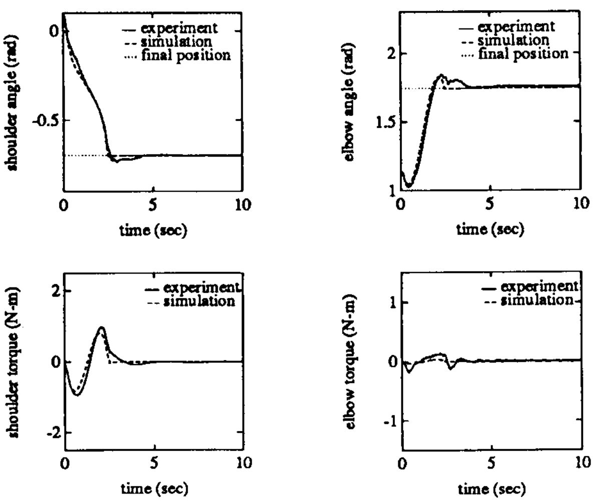

Figure 8: **Terminal Controller Experimental Response for $t_f = 2.5$ seconds**

Comparison of experimental and simulated position and torque histories for the shoulder and elbow motors.

correspond quite well with the simulated results, with the exception that the experimental responses show slightly more

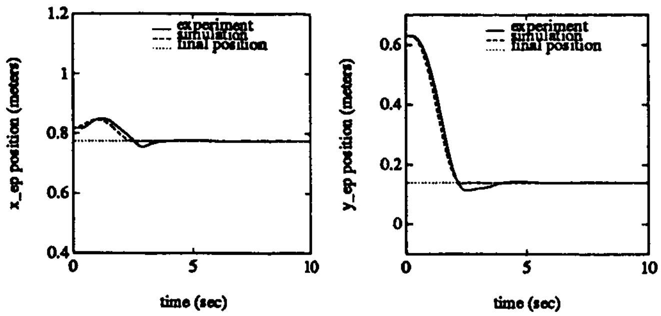

Figure 9: **Experimental End-Point Motion for $t_f = 2.5$ seconds**

Experimental and simulated end-point response for the terminal controller.

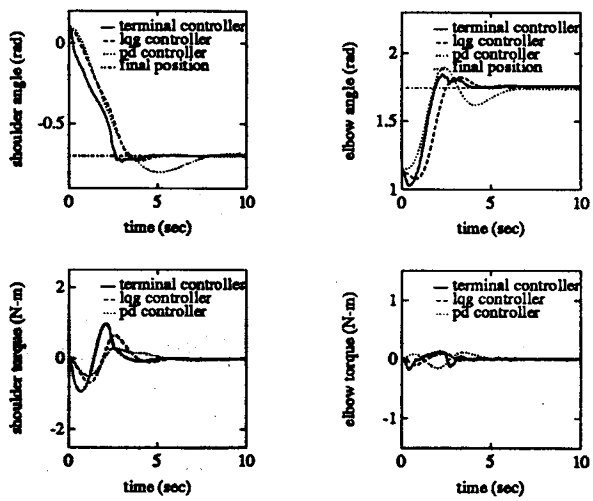

Figure 10: **Controller Comparison for $t_f = 2.5$ seconds**

A comparison of experimental position and torque histories for the terminal controller, an endpoint based LQG controller, and a PD controller.

overshoot in both cases. The endpoint position response shown in Figure 9 again shows excellent agreement between experiment and simulation.

Figure 10 gives a comparison of the terminal controller to an LQG end-point-based controller and a PD controller. The response of the PD controller is by far the worst, as expected, in that it exhibits excessive overshoot and a large settling time. Considering the end-point response, Figure 11 illustrates that the terminal controller shows significantly shorter rise times and settling times than the LQG controller, with comparable overshoot [12].

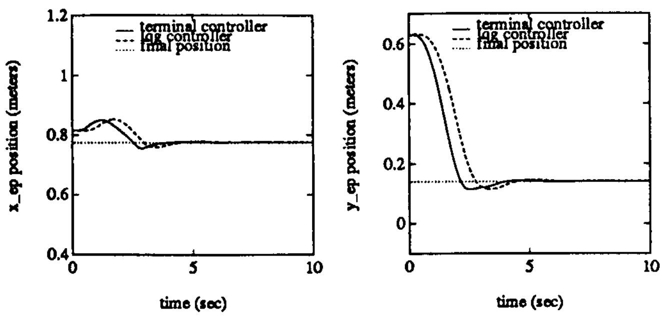

Figure 11: Controller Comparison of End-Point Motion for $t_f = 2.5$ seconds

Experimental and simulated endpoint responses for the terminal controller and an end-point based LQG controller.

Acknowledgements

The research reported here has been funded by the National Aeronautics and Space Administration and the Air Force Office of Scientific Research. The authors would like to thank Professor Robert H. Cannon Jr. and all of the students and technical staff at the Stanford Aerospace Robotics Laboratory for their support, assistance, and suggestions.

References

[1] S. Schneider. *Experiments in the Dynamic and Strategic Control of Cooperating Manipulators.* PhD thesis, Stanford University, Stanford, CA 94305, September 1989. Also published as SUDAAR 586.

[2] M. A. Ullman. *Experiments in Autonomous Navigation and Control of Multi-Manipulator Free-Flying Space Robots.* PhD thesis, Stanford University, Stanford, CA 94305, (February) 1992. To be published.

[3] Vincent W. Chen. *Experiments in Adaptive Control of Multiple Cooperating Manipulators on a Free-Flying Space Robot.* PhD thesis, Stanford University, Stanford, CA 94305, (September) 1992. To be published.

[4] William C. Dickson. *Experiments in Cooperative Object Manipulation by Free-Flying Robot Teams.* PhD thesis, Stanford University, Department of Aeronautics and Astronautics, Stanford, CA 94305, (October) 1992. To be Published.

[5] Eric Schmitz. *Experiments on the End-Point Position Control of a Very Flexible One-Link Manipulator.* PhD thesis, Stanford University, Department of Aeronautics and Astronautics, Stanford, CA 94305, June 1985. Also published as SUDAAR 548.

[6] Celia M. Oakley. *Experiments in Modelling and End-Point Control of Two-Link Flexible Manipulators.* PhD thesis, Stanford University, Department of Mechanical Engineering, Stanford, CA 94305, April 1991.

[7] Robert H. Cannon, Jr. and Eric Schmitz. Initial experiments on the end-point control of a flexible one-link robot. *The International Journal of Robotics Research*, 3(3):62–75, Fall 1984.

[8] Daniel M. Rovner and Robert H. Cannon, Jr. Experiments toward on-line identification and control of a very flexible one-link manipulator. *International Journal of Robotics Research*, 6(4):3–19, Winter 1987.

[9] Daniel M. Rovner and Gene F. Franklin. Experiments in load-adaptive control of a very flexible one-link manipulator. *Automatica*, 24(4):541–548, July 1988.

[10] Lawrence Alder. *Experiments in Adaptive Control of a Flexible Link Robotic Manipulator with Unknown Dynamic Payloads*. PhD thesis, Stanford University, Stanford, CA 94305, (January) 1992. To be published.

[11] L. J. Alder and S. M. Rock. Control of a flexible robotic manipulator with unknown payload dynamics: Initial experiments. In *Proceedings of the ASME Winter Annual Meeting*, pages 61–66, Atlanta GA, December 1991. ASME.

[12] W. L. Ballhaus, S. M. Rock, and A. E. Bryson, Jr. Optimal control of a two-link flexible manipulator using time-varying controller gains: Initial experiments. In *Proceedings of the American Astronomical Society Guidance and Control Conference*, Keystone, CO, February 1992.

[13] C. M. Oakley and R. H. Cannon, Jr. Anatomy of an experimental two-link flexible manipulator under end-point control. In *Proceedings of the IEEE Conference on Decision and Control*, Honolulu, HI, December 1990.