

Mobile Worksystems for Decontamination and Decommissioning Operations

**Final Report
February 1997**

RECEIVED
SEP 18 1997
OSTI

Work Performed Under Contract No.: DE-AC21-92MC29104

For

U.S. Department of Energy
Office of Environmental Management
Office of Technology Development
1000 Independence Avenue
Washington, DC 20585

U.S. Department of Energy
Office of Fossil Energy
Federal Energy Technology Center
Morgantown Site
P.O. Box 880
Morgantown, West Virginia 26507-0880

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

By

Carnegie Mellon University
Office of Sponsored Research
5000 Forbes Avenue
Pittsburgh, Pennsylvania 15213

MASTER

Disclaimer

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

DISCLAIMER

Portions of this document may be illegible electronic image products. Images are produced from the best available original document.

Abstract

This project is an interdisciplinary effort to develop effective mobile worksystems for decontamination and decommissioning (D&D) of facilities within the DOE Nuclear Weapons Complex. These mobile worksystems will be configured to operate within the environmental and logistical constraints of such facilities and to perform a number of work tasks. Our program is designed to produce a mobile worksystem with capabilities and features that are matched to the particular needs of D&D work by evolving the design through a series of technological developments, performance tests and evaluations.

The Phase I effort was based on a robot called the Remote Work Vehicle (RWV) that was previously developed by CMU for use in D&D operations at the Three Mile Island Unit 2 Reactor Building basement. During Phase I of this program, the RWV was rehabilitated and upgraded with contemporary control and user interface technologies and used as a testbed for remote D&D operations. More significantly, we established a close working relationship with the DOE Robotics Technology Development Program (RTDP).

In the second phase, we designed and developed a next generation mobile worksystem, called *Rosie*, and a semi-automatic task space scene analysis system, called *Artisan*, using guidance from RTDP. Both systems are designed to work with and complement other RTDP D&D technologies to execute selective equipment removal scenarios – in which some part of an apparatus is extricated while minimally disturbing the surrounding objects. RTDP has identified selective equipment removal as a timely D&D mission, one that is particularly relevant during the de-activation and de-inventory stages of facility transitioning as a means to reduce the costs and risks associated with subsequent surveillance and monitoring.

In the third phase, we tested and demonstrated core capabilities of *Rosie* and *Artisan*; we also implemented modifications and enhancements that improve their relevance to DOE's facility transitioning mission. To prove that *Rosie* is capable at the task level, we picked a subset of representative D&D operations and executed them in cold tests. Through these exercises we qualified and quantified important operational characteristics of the mobile worksystem. In parallel, we exercised the *Artisan* TSSA system in realistic scenes and modified its underlying algorithm to recognize more complicated objects. *Artisan* was integrated to *Rosie* and exercised in experiments that measured telerobotic task execution times relative to teleoperation. Most significantly, DOE ordered a second copy of *Rosie* ("*Rosie-C*") to participate in one of the D&D Focus Area's Large Scale Demonstrations, dismantlement of the CP-5 test reactor at Argonne National Laboratory.

Disclaimer

The descriptions, accounts and conclusions presented herein are those of Carnegie Mellon University and RedZone Robotics, Inc. and do not necessarily represent those of the US Department of Energy or any other Federal agency.

Acknowledgments

This work was performed under contract No. DE-AC21-92MC29104 for the Department of Energy, Morgantown Energy Technology Center.

We acknowledge and appreciate the interest and constructive support of the DOE technical project officer, Mr. Vijendra P. Kothari, as well as that of Dr. Linton Yarbrough of the DOE EM Robotics Technology Development Program; Mr. Dennis Haley, Dr. William Hamel, Mr. Mark Noakes, Mr. Robert Barry and Dr. John Draper of Oak Ridge National Laboratory's Telerobotics Systems Section; and Mr. Daniel Horschel, Mr. Charles Little, Mr. Michael McDonald and Dr. Michael Griesmeyer of Sandia National Laboratories' Intelligent Systems Group.

Table of Contents

1. Executive Summary	1
1.1 Introduction	1
1.2 Summary Description of Results	4
1.2.1 Rosie Mobile Worksystem	4
1.2.2 Artisan Task Space Scene Analysis System	7
1.3 Guide to the remainder of this report	9
1.4 For more information	9
2. Background	11
2.1 Design of the Original RWV	11
2.2 Decontamination and Decommissioning Needs Analysis	12
3. Results	15
3.1 Rehabilitation of the TMI Remote Work Vehicle	15
3.2 Testing of RWV I in Mock-ups	15
3.3 Rosie Mobile Worksystem - Original Design	18
3.3.1 Locomotor	19
3.3.2 Heavy Manipulator	21
3.3.3 Tooling and Auxiliary Services	22
3.3.4 Control System	23
3.3.5 Operator Console	24
3.3.6 Power and Telemetry	25
3.4 Cold testing of Rosie (Phase III)	25
3.4.1 Testing Program Summary	25
3.4.1.1 Objectives	25
3.4.1.2 General Results	27
3.4.2 General System Characterization	28
3.4.3 Cold Testing of D&D Tasks	30
3.4.3.1 Concrete Demolition	30
3.4.3.2 Metal Cutting	31
3.4.3.3 Debris Removal	32
3.5 Modifications and enhancements of Rosie	32
3.5.1 Hydraulics	32
3.5.2 Locomotor	33
3.5.3 Heavy manipulator	34
3.5.4 Tether Management System	36
3.5.5 Control System	36
3.5.6 Console	37
3.6 Artisan Task Space Scene Analysis System	38
3.6.1 Fundamental Aspects of Artisan	39
3.6.2 Quadric/Planar Segmentation and Matching Method	42

3.6.3 Free-Form Object Recognition Method	43
3.6.3.1 Introduction	43
3.6.3.2 FORM Operating Scenario	44
4. Conclusions	51
4.1 Mobile Worksystem Characteristics	51
4.1.1 Configuration	51
4.1.2 Locomotion	52
4.1.3 Control	52
4.1.4 Manipulation	53
4.1.5 Tooling	54
4.1.6 Operator Interface	55
4.2 Task Space Scene Analysis Systems	55
5. Recommendations	58
5.1 Future Modifications to Rosie	58
5.1.1 Decontamination	58
5.1.2 Operations	58
5.1.3 Alternative Configurations	59
5.2 Future Modifications to Artisan	59
5.2.1 Recognition of Individual Objects	59
5.2.2 Recognition of Multiple Objects	60
5.2.3 Use of Range Data	61
5.2.4 User Interface	61
5.3 Integrated Systems & Demonstrations	62
6. References	64

List of Figures

Figure 1.1	The Rosie mobile worksystem	5
Figure 1.2	Artisan task space scene analysis example	8
Figure 3.1	Rosie with heavy manipulator in home (stowed) configuration	19
Figure 3.2	Rosie steering modes	20
Figure 3.3	Rosie heavy manipulator in several configurations	22
Figure 3.4	Rosie operator console	24
Figure 3.5	Overall Artisan screen layout	40
Figure 3.6	Scanning laser rangefinder	41
Figure 3.7	Range image (left) and reflectance image (right) of experimental testbed	42
Figure 3.8	Cartesian mesh formed from range data	43
Figure 3.9	Range data used in the FORM recognition example	44
Figure 3.10	Cartesian mesh before and after 3-D smoothing filter	45
Figure 3.11	Steps to add a model to Artisan's library	46
Figure 3.12	Spin image basis geometry and example spin images	48
Figure 3.13	Objects recognized using FORM overlaid on Cartesian meshes	49
Figure 3.14	TeleGRIP workcell showing objects recognized using FORM	50

List of Tables

Table 2.1	Overview of D&D tasks & tools	12
Table 2.2	General approach to a D&D mission using robotic technology	14
Table 3.1	Phase I cold testing results	15
Table 3.2	Rosie locomotor specifications	20
Table 3.3	Partial list of tools compatible with Rosie	23
Table 3.4	Artisan software modules	39
Table 3.5	Objects in Artisan's library	46
Table 4.1	Degrees of freedom for several mobile worksystem configurations	53

1. Executive Summary

1.1 Introduction

Part of the Cold War legacy is hundreds of facilities within the US nuclear weapons complex that now sit dormant. Most of them are contaminated radiologically, chemically, or both. At worst, many of them represent a very serious ecological threat; at best, all of them continue to be economic burdens due to necessities of security, surveillance and general maintenance. These considerations motivate decommissioning of the facilities, itself an expensive proposition, but the only means by which this chapter of the Cold War can be closed. Decontamination and dismantlement (D&D) of nuclear facilities is a genuine problem that cannot be ignored and one whose solution is likely to cost billions of dollars. The magnitude alone of the task is daunting; the situation is exacerbated due to lack of technologies to execute many parts of the work agenda.

The intent of the project described in this report was to innovate technical solutions that would enable some aspects of nuclear facility decommissioning to proceed. Further, and in keeping with the Department of Energy's strategy of investing in research and development that ultimately lowers the cost and reduces the time scale of environmental restoration, the aim of the project was to create technologies that would be commercially available both to DOE and its contractors. Specifically, we set out to develop a state-of-art mobile robot worksystem that could deploy itself into a nuclear facility and perform tasks that would otherwise require one or more human workers. In addition to developing robot hardware that could accomplish remote work, our goal was to develop software that would make the efficiency of the remote worksystem rival that of its human counterparts. Achieving those goals meant realization of capabilities to execute a significant fraction of the total D&D agenda with zero worker exposure and without loss of human worker productivity.

Carnegie Mellon University's Robotics Institute, the prime contractor, pursued this work because it represented opportunities both to fill a real technology gap and to make a positive impact on a problem of national importance. The Robotics Institute emphasizes disciplined engineering of systems – even in projects that have significant research content – because it forces consideration of issues that ultimately determine if a solution has relevance beyond the laboratory. We chose RedZone Robotics, a CMU spin-off company that shares our commitment to solving meaningful problems, as a partner in the project. With a commercial entity playing a major role, we felt there was a higher likelihood of the project's results being rapidly transferable into practice. To round out the team, we garnered technical support from the DOE Robotics Technology Development Program, Oak Ridge National Lab's Telerobotic Systems Section and Sandia National Lab; Morgantown Energy Technology Center provided project administration and ties to the D&D Focus Area.

The genesis of the project dates to the mid-80's, when CMU was heavily involved in the decommissioning of Three Mile Island's Unit II reactor building. In that effort, CMU designed and built three mobile robots for recovery of the reactor building basement. The first two of these machines performed reconnaissance missions using remote cameras, radiological sensors, and core drill samplers to help quantify the environmental parameters that eventually defined the scope and strategy of the clean-up. The third robot, the *Remote Work Vehicle* or *RWV*, was developed to carry out the work agenda established by its predecessor machines. RWV was

designed to be a multipurpose machine that combined brute force with finesse to perform a wide variety of work tasks such as sludge scooping, pipe cutting, concrete demolition, surface scabbling, debris packaging, coating application, and others. The RWV was serial chain of remotely operated components: an omni-directional locomotor, attached to which was a high-strength, high-reach boom, attached to which was a spatially correspondent manipulator (for dexterous work) and heavier tools that the manipulator could not carry (such as a 12-inch concrete coring saw). RWV had many design features that made it particularly well-suited for remote work in a radiological environment: primary motions were hydraulic, critical components had backups, finishes allowed aggressive decontamination, and control was straightforward. As a tool for decommissioning, RWV was state-of-art in its time, but despite its capability, however, RWV was relatively simplistic in that it was completely remote controlled and it took three skilled operators to run it.

Ultimately, RWV was not put into service due to programmatic considerations that forced greater attention to be given to other aspects of the overall TMI Unit II decommissioning, and CMU was able to reacquire the RWV. In principle and intent, RWV was still a very viable approach to D&D, hence it was well worth consideration for application in the nuclear weapons complex.

During Phase I of this project (October 1992 - September 1993), the RWV was rehabilitated and exercised as a testbed for remote D&D operations specific to the needs of US DOE. The objective of this first phase of the project was to ascertain what design features of RWV were appropriate and which ones, if any, required modification. Our initial assessment indicated that, from a control standpoint, the multiple operator for a single machine operating scenario was too cumbersome and cost inefficient. RWV was therefore upgraded with contemporary computer control and user interface technologies prior to cold testing. The tests themselves consisted of numerous work tasks representative of D&D operations as well as "generic" actions that described the fundamental capabilities of the combined man-machine system (e.g., driving maneuvers and object grasping and placement). Results of the Phase I testing program proved that many salient aspects of the RWV design – general morphology, multiplicity of purpose, size, reach, power, maneuverability, and speed – were right for D&D. However, many underlying aspects, such as extremes of redundancy, were overly complex. Operation of the worksystem, despite reduction in number of operators from three to one, was also lacking given the need to work productively and cost-effectively. Teleoperation persisted as the only available mode of control; while offering the benefits of reduced human exposure and occupational hazard, teleoperation limits progress by its very nature, sometimes to one-tenth of the equivalent speed of manual (hands-on) labor.

Phase I results indicated that development of new technologies was warranted. First, a new generation of mobile worksystem, called *Rosie* (Figure 1.1), was conceived, engineered, and fabricated. The *Rosie* design retained all of the strengths of the RWV and alleviated many of its shortcomings; key improvements over its predecessor included increased power, greater strength, better control and overall design simplification. The second technology developed was a computer vision system, called *Artisan* (Figure 1.2) to scan and generate a three dimensional model of the space in which a robot performs its task. This paradigm, known as task space scene analysis or TSSA, enabled remote equipment such as *Rosie* to work with some degree of autonomy, thus improving upon the baseline of teleoperation. Development of *Rosie* and *Artisan* was guided by the DOE Robotics Technology Development Program (RTDP), with whom we had established a

strong working relationship, to help ensure end-user relevance as well as technological compatibility and programmatic synergy with other D&D robots and control systems.

In Phase III we tested and demonstrated core capabilities of Rosie and Artisan; we also implemented modifications and enhancements that improve their relevance to DOE's facility transitioning mission. To prove that Rosie is capable at the task level, we picked a subset of representative D&D operations and executed them in cold tests to qualify and quantify important operational characteristics of the mobile worksystem. In parallel, we exercised the Artisan TSSA system in realistic scenes and modified its underlying algorithm to recognize more complicated objects. At the end of Phase III, the two technologies we developed were married together: Artisan was integrated to Rosie and exercised in experiments that measured telerobotic task execution times relative to teleoperation. Rosie is now commercially available; Artisan is functional as a laboratory-scale prototype and requires some additional enhancement prior to commercialization (anticipated in mid-1998).

Phase III testing verified that the technologies we developed are viable and useful for D&D.

- Rosie can capably perform a wide range of tasks and function within the operating constraints of DOE facilities.
- Artisan can quickly and accurately make a three dimensional model of a D&D worksystem's task space.
- Task space scene analysis technologies (like Artisan) and D&D work robots (like Rosie) can be readily integrated.
- Worksystems that use Artisan can work semi-automatically.
- Semi-automatic control can be superior to teleoperation for tasks that are repetitive, involve high precision work, or require extreme operator patience.

Perhaps the most telling result of the program was fulfillment of our principal objective: to elicit a request from a DOE facility to employ our technologies, and this has occurred. DOE ordered a second copy of Rosie ("Rosie-C") to participate in one of the D&D Focus Area's Large Scale Demonstrations, dismantlement of the CP-5 test reactor at Argonne National Laboratory.

Rosie is now commercially available – it is RedZone's flagship product – and additional markets beyond nuclear weapons facility D&D are also being sought. These include construction, aircraft maintenance, commercial nuclear plant decommissioning and other applications that stand to benefit from high-reach, high-payload mobile manipulators that can work in modes ranging from teleoperation to telerobotic to autonomous. We estimate that additional copies of Rosie will cost between \$750k and \$1.25M, depending on particular task requirements.

Artisan has not yet reached the same maturity level as Rosie and additional work of three kinds is needed before it is ready for field use. First, some technological improvements can and should be made that extend its utility. These include expanding the library of objects that can be recognized, making it work with data sets that are larger than a single range image, and giving it the ability to scale recognition results to increase fit accuracy. Further, through code optimization and faster processing, it is possible to make the entire recognition process almost completely automatic. An operator would then only have to point the range sensor, view the results and confirm them – no other interaction would be needed. Second, the entire system needs to be brought up to a

commercial standard, which will at least entail some additional, more rigorous software engineering, and should also include porting the system to less expensive computational platforms, e.g., PCs.

Finally, there is substantial effort ahead to convince those who are actually responsible for D&D jobs to put Rosie and/or Artisan into practice. To the non-roboticist, our technologies are often viewed more as technical curiosities than as tools, and this perspective can only be changed by being truly responsive to needs and wants of the end-user community. While we argue that robots will save money in the long run, our contentions are rather difficult to defend given our price tags and the lack of real cost/benefit numbers. Our technologies will ultimately justify themselves, but gaining entry into the whole D&D arena is a necessary first step, and additional nurturing is required beyond that. It is therefore critical to push on demonstration activities that will generate hard evidence in fair tests against existing alternatives. There must also be some latitude afforded to D&D job managers that will alleviate performance, cost and schedule pressures to an extent that encourages them to give these technologies the opportunity to succeed.

1.2 Summary Description of Results

1.2.1 Rosie Mobile Worksystem

The Rosie mobile worksystem provides all of the necessary locomotion, heavy lifting and tooling deployment capabilities to perform D&D tasks remotely. Surplus hydraulic power, electrical capacity and control channels allow the addition of a wide variety of tools ranging from hand-held devices to large demolition and decontamination equipment. Rosie is an electrohydraulic omnidirectional locomotor with a four degree of freedom heavy manipulator mounted on its deck. The heavy manipulator has a payload capacity of over 2000 lbs throughout a work envelope that extends from 0 to 30 ft. above floor level. Rosie can deploy large tools for demolition and decontamination or carry smaller manipulators for more dexterous work using remotely operated hand- and power-tools. Rosie is designed to withstand the rigors of heavy work over periods of years and is amenable to decontamination so that it can be safely transported from worksite to worksite. Rosie's coordinated controls, ergonomic user interface, and modular design simplify installation and operation of the worksystem.

The locomotor itself is four-wheel drive/four-wheel steer mechanism that provides and distributes hydraulic and electrical power and computing capabilities for all manipulation, tooling and sensing functions. A multi-conductor umbilical managed by an on-board reeling system connects the robot to an off-board electrical power distribution unit and control console. The heavy manipulator is a high-payload, telescoping boom that can carry up to 2000 lbs at full extension.

All principal Rosie motions are servo controlled by a real-time computer system. all movements can be directed by a human operator at a remote console or by another computer. This affords a variety of control modes including joint-level teleoperation operator-directed movement of several computer-coordinated axes, teach/playback of short duration heavy manipulator motion sequences, and fully automatic execution of work tasks under the direction of an external computer.




Figure 1.1 The Rosie mobile worksystem

Rosie's control and communications architecture is designed to accommodate the addition of sensors connection to additional processors through RTDP- and industry-standard interfaces. Multiple on-board cameras, with pointing and lens control functions, provide remote viewing capabilities, and stereo microphones provide audio feedback.

To assure programmatic fit in the overall RTDP, Rosie was designed with substantial consultation from the D&D Robotics Coordinator and technical experts from ORNL. To assure sound

engineering design, additional expertise was employed for finite element analyses of structural components, electrical design in accordance with ORNL standards, console layout, software design, and hydraulic system analysis. Rigorous quality assurance practices were followed throughout the component procurement and fabrication process.

The Rosie design incorporates many onboard and off-board features that facilitate remote operation. Rosie's four wheel drive/steer configuration makes the locomotor highly maneuverable in tight or cluttered spaces. The front axle extends outward for added stability when working at elevation and the lockable pivot-mounted rear axle provides compliance when working on uneven floors and crossing obstacles. Critical driving and steering motions are functionally redundant — each wheel module is individually driven, steered and has surplus power to compensate for limited failures. The wheel drive motions free-wheel when unpowered to enable emergency recovery towing.

Reliability is essential in environments where manual recovery of failed equipment is difficult, costly, or precluded by hazards. The rugged construction of this system is suited to the abusive conditions of dismantlement operations and it is designed to withstand inadvertent collisions or falling objects. Rosie is tethered for reliable communications and power delivery in extended work periods. Onboard tether management ensures that the tether is not endangered by being dragged. The communication system has high noise immunity and error-checking to ensure faithful transmission of control signals between the operator's control console and the robot. A suite of on-board sensors monitors the status of critical components and automatically alert the operator of potential problems.

The robot portion of the system is designed to withstand a cumulative radiation dose of 10^5 R. Materials and components have been selected to reduce the potential for radiation degradation. Higher levels of radiation hardening are achievable if necessary by shielding of critical electronics and using more radiation tolerant components. Since the ability to decontaminate Rosie (to allow maintenance, storage, and transportation to other facilities) with a minimum of personnel exposure is critical, all onboard components are sealed for pressurized washdown. The robot's structures are designed to minimize exposed surfaces and areas where contamination can collect and be trapped and areas that can't be sealed are as open as possible to facilitate cleaning and washdown.

Automation of low-level functions and other control system features allow a single operator to maneuver both the locomotor and heavy manipulator and to work very efficiently. All motions incorporate position sensing and servo-control, enabling precise motion control for dexterous positioning either by teleoperation or computer control. High resolution and continuously variable speeds allow an operator to move slowly for fine positioning, or quickly for efficient large motions. Up to 10 onboard video cameras — with lights and pan/tilts — provide an operator with effective views for navigation and tool deployment.

Rosie is designed to be as modular as possible to expedite maintenance and deployment of alternate tools. Modularity allows the quick replacement of components or subsystems in order to keep the system in service, and allowing failed components to be repaired off-board and off-line. Critical components in the electrical and hydraulics systems are readily accessible and can be modularly replaced.

1.2.2 Artisan Task Space Scene Analysis System

Remote operation offers the advantage of reduced human exposure to hazardous working conditions, but there is a substantial penalty paid in reduced overall speed and productivity. Much of the efficiency loss is attributable to the relatively low bandwidth of teleoperation man-machine interfaces; the human operator is responsible for planning, coordinating and sequencing actions, though the only information available to him is gathered by remote sensors and presented on 2-D displays. Productivity of a mobile worksystem can be increased if it can function with less dependence on a human. Automatic operation is exactly what has been achieved with assembly line robots; their cost-effectiveness relative to manual labor is undisputed. By extending this principle into D&D, i.e., if some amount of autonomy can be imbued on D&D worksystems, we can expect a similar gain in productivity.

The key to automation is knowledge of geometry: of the objects the robot deals with and how they are arranged in the robot's workspace. Early factory robots operated in very structured environments, that is, the size, shape and location of objects relative to the robot were known ahead of time. So it was possible to pre-program robot actions that were virtually guaranteed to succeed as long as the *a priori* information remained valid. This required rigging the environment of the robot such that the geometric information implicit in the robot program remained correct. Modern factory robots are more capable because they are equipped with sensors that allow them to make measurements and recognize objects during program execution. Greater flexibility is afforded since the robot can compensate for geometric variations and react to changing circumstances. In the realm of D&D, however, the environment is poorly structured and by definition very dynamic. Robotic D&D work tasks therefore defy pre-programming and require more sophisticated means to acquire correct information on the fly.

The Robotics Institute has deep history in computer vision, thus we were well suited to develop technology that would complement Rosie and other D&D worksystems by making them more automatic. With guidance from RTDP, we developed Artisan, a perception system that combines range sensors and object analysis software to create a 3-D model of a robot's work area – *in situ*, i.e., as it is discovered by the robot. This paradigm, known as task space scene analysis, provides the information necessary for a computer to calculate safe robot motion trajectories and tooling paths.

While there are alternative approaches to task space scene analysis currently under development, Artisan offers distinct advantages. Scene analysis techniques that map space into occupied and unoccupied regions would be sufficient to plan collision-free motion trajectories. Artisan goes beyond this, however, to provide additional semantic information such as function, likely material type and even manufacturer, is made available along with shape, size and location. This is because Artisan actually recognizes objects which allows both the robot and its operator to understand the not just the layout, but also the context of the work space and the objects in it.

A second benefit offered by Artisan is its ability to recognize virtually any object for which there is an existing CAD model. And for objects for which no such data exists, the underlying object recognition algorithm allows objects that have been seen and measured once to be recognized in the future. Finally, and perhaps most significantly, Artisan itself is semi-automatic and requires very little input from its operator to function. Artisan thus creates the opportunity to elevate operator interaction with D&D worksystems from man-in-the-loop teleoperation to supervisory/

teleoperated control and bring the benefits of safeguarding, finesse, repeatability and productivity of automation to bear on tasks that have been traditionally relegated to remote control.

World modeling with Artisan proceeds as follows. From the console, an operator positions a range measurement device (such as a scanning laser rangefinder) and acquires range data that are displayed as images on his computer screen. The operator then defines a region of interest within the display and selects objects that appear in it from a CAD model database. From these clues, Artisan tries to recognize the objects in the scene, and if it finds them, calculates their position and orientation. The operator then has the opportunity to accept or reject those results. If a recognition result is accepted, the object(s) are added to a virtual world model that also includes the worksystem. This process of range data collection, processing, and user confirmation continues until all desired objects pertinent to the task at hand have been recognized. With the 3-D task space model in place, the operator can then access automation features such as trajectory planning, collision avoidance, and scripted motion sequences to execute tasks.

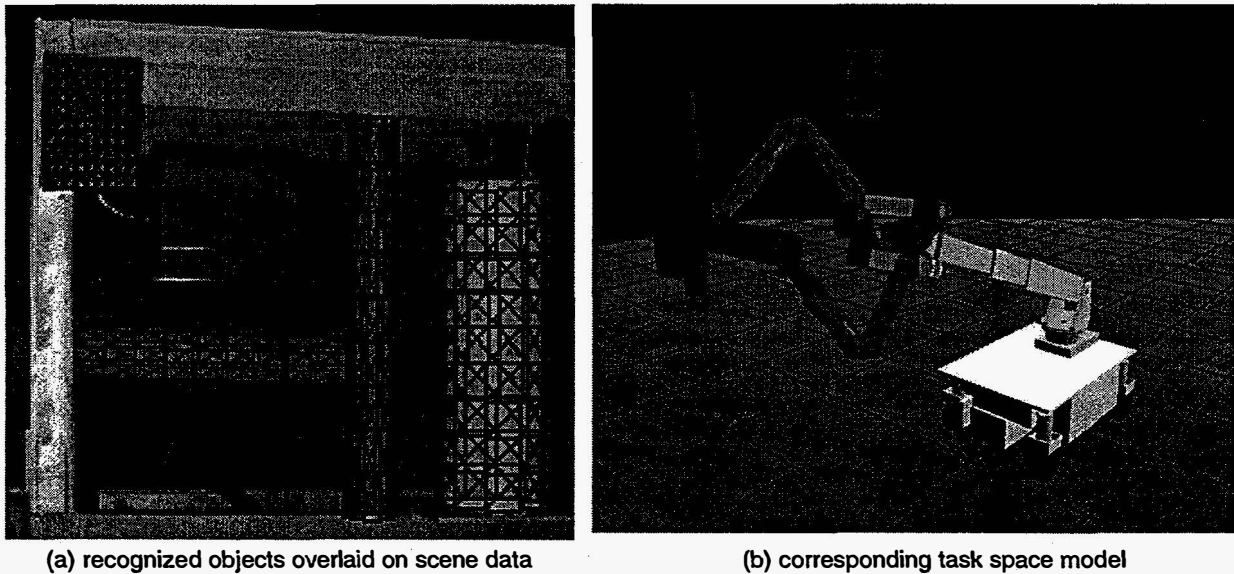


Figure 1.2 Artisan task space scene analysis example

Two different object recognition algorithms have been developed for Artisan. The first method (Quadric/Planar Segmentation and Matching, or "QPSM") segments the 3-D surfaces into planar and quadric patches and matches the resulting scene description to analogous descriptions of object models in a database (developed off-line from CAD descriptions of objects). The other method (Free-form Object Recognition Method or "FORM") is based on a technique known as geometric indexing. In this case each point in a 3-D collection of points is transformed into a 2-D representation, called a spin image, that describes the spatial relationship of all the other points to the selected point. Each point in a 3-D data set – whether from a range sensor or from a CAD model – has a corresponding spin image. The first pass in the recognition algorithm is comparison of random samples of spin images representing the scene data to spin images of models in the object database, which results in a few plausible correspondences between data and library model. In the second step, plausibility is further tested and the best matches are refined using an

algorithm that outputs an optimal estimate of the recognized object's location and orientation in the task space.

Artisan exports its results into an interactive robot simulation package, TeleGRIP™ by Deneb Robotics, Inc. TeleGRIP keeps track of the geometry of the entire task space including all recognized objects, all *a priori* information and the remote worksystem itself. TeleGRIP provides visualization capabilities that allow the operator to view the task space and because the model is synthetic, the operator can change the perspective to see behind and underneath objects even if it is impossible to move a real camera into that position. This process of range data collection, processing and user interaction continues until the operator is satisfied with the 3-D model of the robot's work space. A TeleGRIP workcell constructed using Artisan (also showing Rosie carrying a DOE two-arm manipulator) is shown in Figure 1.2b.

More importantly, the operator can command the robot to perform a task by running a simulation. TeleGRIP provides several intrinsic motion planning capabilities as well as hooks that allow additional task programming scripts (developed off-line) to be executed at run time. For example, in a decontamination operation, the operator would like the robot to follow the surface of a pipe while maintaining a particular stand off distance. The operator selects a starting point for the end effector and the appropriate trajectory is calculated automatically. During simulation of task execution, TeleGRIP will warn the operator of any collisions or motions that exceed joint limits of the worksystem. Once satisfactory performance is obtained, the operator can "download" the resulting tool paths to the real worksystem.

1.3 Guide to the remainder of this report

Background information, including the original RWV design and DOE's needs for remote D&D worksystems, are provided in Section 2.

Results of the program are found in Section 3, which is subdivided into subsections on

- rehabilitation of the RWV (Section 3.1),
- Phase I testing of the RWV in mock-ups (Section 3.2),
- a more detailed description of Rosie (Section 3.3),
- cold testing and characterization of Rosie (Section 3.4),
- enhancements and modifications of the original Rosie design (Section 3.5), and
- a more detailed description of Artisan (Section 3.6).

Conclusions are presented in Section 4, recommendations are presented in Section 5, and references in Section 6.

1.4 For more information

While we have attempted to convey complete descriptions of Rosie and Artisan in this document, fuller appreciation of the technologies is possible. Two prior reports describing this project have been published by DOE and are available from DOE (Office of Scientific and Technical Information, PO Box 62, Oak Ridge, TN 37381) or the National Technical Information Service (Springfield, VA 22161). The earlier report (DOE document #DOE/MC/29104-3662 or NTIS document DE94004079) describes rehabilitation of RWV and cold testing performed in Phase I,

while the second (DOE document #DOE/MC/29104-5296 or NTIS DE97002034) describes the Rosie mobile worksystem in detail and describes the earlier version of Artisan based on quadric/planar segmentation and matching.

Promotional videos describing Rosie are available from RedZone Robotics. Please contact:

Leona Bares
RedZone Robotics, Inc.
2425 Liberty Avenue
Pittsburgh, PA 15222
phn: 412-765-3064
fax: 412-765-3069
email: lbares@redzone.com

A video describing Artisan is due for completion in late February, 1997, and a CD-ROM with a demo version of the software will be available in late March, 1997. Please contact:

Jim "Oz" Osborn
The Robotics Institute
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15213
phn: 412-268-6553
email: oz@cmu.edu

The web site for this project, which includes links to some downloadable publications, is:

<http://www.frc.ri.cmu.edu/rww/Artisan.html>

2. Background

2.1 Design of the Original RWV

The Remote Work Vehicle (RWV) is one of three mobile teleoperated systems developed for accident recovery activities at Three Mile Island Unit-2. It was conceived to fulfill a specific set of task objectives and had unique attributes for operation in the Reactor Building Basement (RBB) where radiation exposure precluded the use of a human work force. The RWV design was based on the experience of predecessor robots at TMI-2 and technologies, system philosophies and operating methods proven to work in other adverse environments.

At the time of the TMI-2 accident, existing remote worksystems for nuclear applications were poorly matched to the work scope and operating conditions for RBB recovery. GPU Nuclear, Inc., thus commissioned the development of a new remote system, the RWV, for recovery efforts in the RBB. In addition to access and maneuvering requirements, performance capabilities specified for the RWV included the ability to deploy tooling from floor to ceiling (22 feet high) for tasks including de-watering and sediment removal, concrete demolition, steel structure dismantlement, surface decontamination and coating application, and material transport. Goals were also set for reliability (assured egress), operability (remotely, by a three person team), decontaminability (wash down with high pressure borated water), and extensibility (for additional tasks).

The RWV's design philosophy was influenced by operating experience and component technology of its predecessor TMI-2 reconnaissance robots, principally the Remote Reconnaissance Vehicle (RRV) and Remote Core Sampler (RCS). Where possible, RWV development sought pre-existing hardware solutions and successful system philosophies from other industries to expedite the development and assure performance. Many conditions in the RBB were similar to the environments in which undersea robots and construction equipment operate routinely. The undersea environment is dark and wet, operators are very remote from the equipment, and successful operation is economically justified despite formidable difficulties. Undersea operations demand reliable equipment over prolonged missions and achieve such objectives through overrating components and effective sealing against water intrusion. Like many nuclear applications, operation and maintenance of undersea remote equipment are performed in the field and systems must therefore be understandable, diagnosable and repairable by field technicians. Since construction equipment must survive nearly constant exposure to dust, dirt and water along with occasions of falling debris or collisions, construction equipment endurance and ruggedness standards were considered appropriate to the needs of long duration recovery operations in the TMI-2 RBB.

To achieve the required work capability and environmental hardening, the RWV used electrohydraulics as its main power source. As vehicle and tooling power requirements increase, electrohydraulic systems begin to have significant size and weight advantages over electric systems. Hydraulic pumps are central power supplies that can support the numerous motions required of capable work equipment and all system power can be focused on a single motion where appropriate. Hydraulic actuators are smaller than electric counterparts of equivalent power, do not typically require gearing (allowing direct-drive and back-drive) and are forgiving in overload and impact tolerance. Many work tools are hydraulic because of the simple conversion of hydraulic pressure to mechanical force.

The reach, stability, and mobility exhibited by construction equipment were also influential. The RWV's telescoping boom takes precedence from high capacity, high reach construction cranes, as does the provision of stifflegs for added stability. The RWV's articulated steering reflects techniques used to achieve mobility on heavy equipment such as rough-terrain cranes, where machine size with respect to the site makes tight maneuvering essential.

2.2 Decontamination and Decommissioning Needs Analysis

The environment in which mobile worksystems will have to perform D&D operations range from areas in which no special clothing is needed to protect workers to areas in which the risk of human entry is great. Dangers can include exposure to alpha, beta and gamma radiation; uranium, plutonium, and tritium; volatile organics; acids and caustics; mercury; TRU waste and mixed waste, and asbestos. Facilities in which D&D is likely to occur include uranium enrichment facilities, including gaseous diffusion plants, centrifuge plants and other separation plants; research and production reactors; hot cells, canyons, and vaults; stacks and cooling towers; silos and waste storage tanks; analytical research labs; and weapons production and assembly facilities. Specific requirements for D&D at DOE defense nuclear facilities can be found in several reference sources; a short summary is presented in our Phase I report DOE/MC/29104-3662.

A partial list of tasks that are candidates for being performed remotely by mobile worksystems is presented in Table 2.1. Given the wide range of tasks that must be executed, the weak structure of the settings in which they are to occur, and the diversity of environmental hazards within the DOE, it is impossible to build one machine to meet all needs. It is more likely that a variety of different systems will be required. These worksystems require versatility to handle a variety of tools and perform a wide range of tasks, they must combine brute force for heavy work with dexterity for fine manipulation, they must be reliable for extended use, and they must be adaptable to a range of work conditions and settings to meet this challenge.

Table 2.1 Overview of D&D tasks & tools

	Tasks	Targets	Tools
Decontamination	<ul style="list-style-type: none"> • Surface washdown • Hot spot removal • Surface removal • Sealant/coating application 	<ul style="list-style-type: none"> • Concrete walls & floors • Structural supports 	<ul style="list-style-type: none"> • pressurized water sprayer • CO₂ pellet blaster • microwave scabblers • laser ablation system • mechanical scabblers • coating sprayer
Component Removal & Size Reduction	<ul style="list-style-type: none"> • Mechanical disassembly • Drilling • Metal plate cutting • Pipe cutting • Crushing 	<ul style="list-style-type: none"> • Ventilation ducts • Electrical power components • Process tanks, vessels & piping • Waste storage drums • Lighting fixtures & supports • Fuel elements 	<ul style="list-style-type: none"> • hydraulic shear • abrasive grinder/saw • reciprocating saw • grapples • drum handler • impact wrench • auxiliary hoist • electromagnet • concrete crusher

Table 2.1 Overview of D&D tasks & tools

	Tasks	Targets	Tools
Demolition	<ul style="list-style-type: none"> • Metal cutting • Concrete cutting • Coring • Drilling • Hammering 	<ul style="list-style-type: none"> • Carbon steel support members & floor grates • Reinforced concrete • Stainless steel reactor vessels, process tanks piping & fuel pool liners • Lead shielding • Concrete floors & walls 	<ul style="list-style-type: none"> • high pressure water • nibbler • hydraulic shear • plasma arc cutter • diamond wire saw • jack hammer • coring saw • pry bar
Debris Collection & Waste Handling	<ul style="list-style-type: none"> • Grasping • Vacuuming • Scooping/shoveling • Waste segregation • Compaction • Containerization • Container transport 	<ul style="list-style-type: none"> • Dry fines • Liquids • Mixed liquids & solids • Concrete rubble • Metal scrap • Nuclear material 	<ul style="list-style-type: none"> • gripper • scoop • bucket • vacuum • sluicer • electromagnet • drum handler
Local Characterization	<ul style="list-style-type: none"> • Sensor deployment • Smear sampling • Sample acquisition 	<ul style="list-style-type: none"> • Process vessel surfaces • Pipe surfaces • Support structures • Walls & floors 	<ul style="list-style-type: none"> • radiation detector • GC/MS probe • specific gas analyzer • surface swipes • coring drill

Equally important are the interfaces and telerobotic controllers that allow human operators to make effective use of remote worksystems to execute a broad work agenda, cope with unforeseen circumstances and react to changing situations. Since it is highly probable that D&D scenarios will always require a man-in-the-loop, the utility of the remote equipment will be largely dependent upon how well the man and machine are able to operate as combined system. This ability to work cooperatively will be determined by the fidelity of interfaces between the worksystems and human operators and the extent to which telerobotic controllers can assume responsibility for safe execution of tasks.

A mobile worksystem is intended to be a useful, usable tool that can be adapted to site needs and quickly put into practice. Modularity of components, both hardware and software, is essential since task requirements, infrastructures, and operating constraints vary across the complex and will change with time as facilities move toward transition. At this point in time, specific needs for D&D are in many respects poorly defined, regulations and policies are in flux, standards for "how clean is clean?" do not yet exist, and overall strategies for meeting D&D requirements in the face of budget limitations are lacking. Furthermore, experience in prior D&D operations is limited, making it difficult to plan, cost, and schedule D&D activities that will span the DOE complex for

the next two decades. Still, it is possible to frame the general work flow of a D&D operation that uses robotic technologies, as shown in Table 2.2.

Table 2.2 General approach to a D&D mission using robotic technology

Task	Sub-tasks
Acquire existing facility information	<ul style="list-style-type: none"> • acquire blueprints, photographs and as-built drawings • assess infrastructure availability & constraints • acquire known contamination data
Acquire existing D&D plans (if available)	<ul style="list-style-type: none"> • interact with EM-30, -40 & -60 • develop understanding of existing D&D plans
Secure facility for D&D mission	<ul style="list-style-type: none"> • select facility • acquire requisite permitting & approvals
Design overall D&D mission	<ul style="list-style-type: none"> • determine D&D end-points • formulate task & tooling requirements • develop computer models of facility • test mission in simulation
Detail D&D tasks	<ul style="list-style-type: none"> • determine overall task sequence • sort tasks by performance with teleoperation or under supervisory control • develop scripts for tasks using teleoperation • develop scripts for tasks using scene analysis & telerobotic control • develop contingency plans to guarantee mission completion
Detail recovery of robots from facility	<ul style="list-style-type: none"> • develop contingency plans for robot failures • develop robot decontamination plan
Prepare robots, sensors and software for mission	<ul style="list-style-type: none"> • develop interfaces between robots, sensors and tooling • develop interfaces between robot controllers & scene analysis sensors • exercise & debug robots, sensors and software • make necessary modifications to robots, sensors & software to accommodate facility & mission
Detail installation of robot equipment in facility	<ul style="list-style-type: none"> • formulate hardware & software installation requirements and develop installation plan • assess installation plan w.r.t. infrastructure availability & constraints
Install robot equipment in facility	<ul style="list-style-type: none"> • interface robots, computational hardware and data communications to facility • verify system & subsystem operation to be within acceptable parameters
Execute tasks	<ul style="list-style-type: none"> • execute tasks using teleoperation • execute tasks using scene analysis & telerobotic control
Recover robots	<ul style="list-style-type: none"> • decontaminate robots • determine corrective maintenance needs • perform routine and corrective maintenance

3. Results

3.1 Rehabilitation of the TMI Remote Work Vehicle

For testing in Phase I, an existing mobile worksystem, the Remote Work Vehicle (RWV) – originally built for use in the Three Mile Island clean-up – was rehabilitated. The RWV testbed is a telerobotic, electrohydraulically powered mobile worksystem equipped with tooling to perform D&D tasks, including decontaminating and sealing concrete and steel structures and surfaces, dismantling equipment, cutting pipes, packaging and transporting materials, acquiring data for radiological and contaminant mapping, and performing general inspections. The need to productively and cost effectively perform such tasks in DOE facilities drove the upgrades to and reconfiguration of the original, purely teleoperated RWV. These upgrades included a new on-board real-time computer control system to coordinate robot motions and monitor status and system health, a new off-board console for one-person remote operation of the robot, and a new suite of tools including a two-axis tool positioner for attachment of heavy tools to the boom tip. A detailed description of the RWV can be found in our Phase I topical report (see the Executive Summary section for the DOE document number).

3.2 Testing of RWV I in Mock-ups

The Phase I testing program was designed to produce qualitative data on the effectiveness of the worksystem testbed for D&D operations. Testing results will be used to develop a better understanding of the capabilities of the RWV relative to needs, requirements and constraints of D&D within DOE facilities and to guide the design of Rosie in the second phase of the program.

We constructed a testing facility that is representative of DOE facilities in which mobile worksystems will operate in the future. The workspace was an industrial setting containing structures, equipment, piping, valving, and ductwork found in processing and assembly plants. The mock-up was designed to be easily reconfigured to test driving and manipulation in several operating scenarios and to allow safe operation of the robot. The methods used in the Phase I testing program are detailed in the Phase I topical report; the description, objectives and results of those tests are presented in the following table.

Table 3.1 Phase I Cold Testing Results

Task	Description	Objectives	Results
Object Recovery	Collect three objects in numerical order (numbered paper targets mounted on the wall) and place them in a container. Three extra targets will be put up as decoys.	<ul style="list-style-type: none"> •Assess difficulty of grasping and placing objects. •Assess adequacy of camera perspectives for object grasping and placing. •Study relative uses of locomotion, boom functions, and manipulation in a task that requires all three. 	<ul style="list-style-type: none"> •Significant time was spent aligning the robot with the target prior to grasping. Set up configuration was lost when the vehicle had to be moved to place objects into a container. •Grasping was not difficult. Judging distance was difficult. •The slew rates of the main camera pan & tilt actuators was too low. The wrist camera was used often since it could be moved rapidly.

Table 3.1 Phase I Cold Testing Results

Task	Description	Objectives	Results
Slalom Driving	Drive through a room, maneuvering around obstacles placed on the floor.	<ul style="list-style-type: none"> •Assess adequacy of on-board camera perspectives in avoiding collisions while driving. •Assess utility of translational and curvilinear driving modes. 	<ul style="list-style-type: none"> •Judging side clearances was perhaps the most difficult aspect of driving. •The manipulator wrist camera was very useful during driving; the main cameras were also useful but they move too slowly. •Translational driving mode was extremely useful; most operators switched between driving modes often.
Pipe Grinding	Given a diagram of pipes hung horizontally on a wall and a specific target pipe, cut through that pipe using a 9-inch grinding disk.	<ul style="list-style-type: none"> •Assess remote operation of the grinder. •Determine limitations of ability to cut closely spaced pipes. 	<ul style="list-style-type: none"> •Grinding was facilitated by locking out the manipulator's waist motion. This helped to prevent binding and assisted in returning to the same cutting groove. •Force feedback would have improved performance in this operation. •The main cameras, which have zoom lenses, were useful in observing work progress. •Use of a foot pedal for on/off control of the grinder worked well.
Pipe Grinding and Recovery	Drive into a room, find the "hot" pipe (identified by a paint marking), cut it down with the grinder and place it into the waste barrel.	<ul style="list-style-type: none"> •Assess system performance while executing a combined driving, dismantling, and debris recovery operation. 	<ul style="list-style-type: none"> •Boom motions were preferred over driving motions when positioning the manipulator prior to cutting. •Picking up the cut pipe sections was difficult. Placing a pipe into a container was relatively easy.
Sawzall Cutting	Cut through a piece of pipe using a reciprocal saw carried by the manipulator.	<ul style="list-style-type: none"> •Assess remote operation of the reciprocal saw. 	<ul style="list-style-type: none"> •The reciprocal saw was best used for soft materials such as wood and PVC. •Keeping the blade in the cutting groove was difficult (several blades were broken during this operation). A mechanism to hold the saw against the pipe would improve performance. •Use of the foot pedal for on/off control of the saw worked well.
Painting	Find a control panel designated on a drawing of the room and paint it, completely covering all surfaces.	<ul style="list-style-type: none"> •Assess remote operation of the paint sprayer. •Assess motions needed for painting surfaces in several different orientations. 	<ul style="list-style-type: none"> •Orienting the robot was far more difficult than actual painting. •The robot's size made it difficult to achieve some orientations/ configurations, particularly for painting downward-facing surfaces.

Table 3.1 Phase I Cold Testing Results

Task	Description	Objectives	Results
Pick Up Sticks	Ten sections of pipe are randomly strewn on the floor of two rooms. Pick up and place into a container as many pipes sections as possible in one hour.	<ul style="list-style-type: none"> •Assess system performance while executing an involved driving and manipulation task. •Force the operator to devise a strategy for performing a task quickly. 	<ul style="list-style-type: none"> •Maneuvering between rooms was difficult. •The main cameras were more useful in assessing the location of all the objects, though the manipulator's wrist camera provided the most useful perspective for picking them up. •Some operators carried the container with them, reducing the time spent to retrieve the pipes. •If the container was upset a great deal of time, up to 10 minutes, was lost righting it.
Pipe Following	Locate a "hot" pipe and move along its length, touching the pipe at one-foot intervals.	<ul style="list-style-type: none"> •Assess ability to collect sensor data in a task analogous to radiological mapping. •Assess ability to combine locomotion and manipulation to follow a long, straight trajectory and judge spatial dimensions. 	<ul style="list-style-type: none"> •Once the robot was aligned perpendicular to the pipe, use of translational driving mode to move sideways made this task straightforward. •Determining distance travelled was difficult.
Control Task	Drive down a hallway, turn a corner, down another hallway, turn into a room through a man door, maneuver to a pipe, and cut down a marked piece with the shears.	<ul style="list-style-type: none"> •Evaluate the operator's ability to control all motions of the RWV. •Compare each operator's performance over multiple runs of the same task and to determine the general shape of the learning curve. 	<ul style="list-style-type: none"> •All but one operator reduced his/her overall time to complete this task by ten minutes during the second week and to half their original time in the third week. •Operators began to fatigue after between 40 and 60 minutes of continuous running. •Repetition of the same task allowed operators to optimize their strategies to reduce execution time.
Valve Manipulation	Locate and shut off two valves. One of the valves is a lever action gas valve, and the other is a rotational water valve.	<ul style="list-style-type: none"> •Assess ability to constrain forces while manipulating objects. •Assess motions needed for manipulating objects where orientations are constrained. 	<ul style="list-style-type: none"> •Controlling forces without any force feedback is very difficult. One circular valve handle was sheared off and one lever action valve was turned past its mechanical stop. •Positioning the robot was a major part of the task.
Pipe Cutting with Two Tools	Maneuver into a room, find a marked pipe and cut it, once using the shears and again using the grinder.	<ul style="list-style-type: none"> •Compare remote operation of the grinder and the shears for pipe cutting. •Evaluate the use of a "tool holster" as part of the Badger heavy tool deployment mechanism. 	<ul style="list-style-type: none"> •The shears are easier to operate but harder to align, since they are deployed by the Badger (the grinder can be positioned with the manipulator). •Relative to the grinder, the shears can cut pipes that are more closely spaced. •Once positioned the shears cut pipes up to four times faster than the grinder.

3.3 Rosie Mobile Worksystem - Original Design

Rosie is a mobile robot worksystem developed for nuclear facility decommissioning and dismantlement. Its primary function is to perform a variety of dismantlement tasks remotely by deploying tools, sensors, and/or other robotic equipment into hazardous areas. Rosie's capabilities and system design address the need for durability and reliability in these environments, and enable performance of tasks such as piping and process equipment removal, structural demolition, vessel segmentation, waste handling and transport, and wall/floor decontamination.

The system includes a tethered robot, a power distribution unit (PDU), and a control console for robot operation. The robot consists of two major subassemblies, the locomotor and the heavy manipulator. The locomotor is a hydraulically powered, omni-directional platform with onboard tether management. It provides mobility to transport the heavy manipulator, tools, or other payloads within the work area. The heavy manipulator is a four degree of freedom, high-payload, long-reach mechanism capable of carrying a variety of tools, one or more dexterous manipulators, or any other payload of up to 2,000 lb throughout a generous work envelope. Rosie is a teleoperated system with low-level automation features that facilitate more efficient remote operations and allow a single operator to maneuver and work effectively.

The overall dimensions of Rosie are 6.5 feet wide by 14 feet long and 8 feet high; the locomotor alone (i.e., without the heavy manipulator attached to its deck) is 6.5 feet wide by 9.5 feet long and 3.5 feet high. The boom telescopes outward; at full extension with the boom horizontal, the overall worksystem length is 19 ft.; with the boom vertical, the reach above the floor is 26.5 ft. Figure 3.1 shows Rosie with the heavy manipulator stowed; Figure 1.1 shows the worksystem with the heavy manipulator near the apex of its work envelope. The overall weight of the machine is about 14,000 lbs.

Rosie's design is modular in nature to allow components to be added or removed, yielding several different configurations. The two main modules are the locomotor and the heavy manipulator. The locomotor provides mobility to carry the heavy manipulator to the work area; the heavy manipulator provides long reach and strength to deploy tools or smaller more dexterous manipulators from its tip. Under this modular concept, the heavy manipulator can be removed from the locomotor allowing other equipment to be deployed from its deck. Concepts include systems such as a variable geometry truss (being developed at Pacific National Laboratory), a scissors lift mechanism, multiple smaller manipulators, floor decontamination mechanisms, etc. Alternatively, the heavy manipulator can be deployed from some other platform, such as an overhead gantry crane.




Figure 3.1 Rosie with heavy manipulator in home (stowed) configuration

The heavy manipulator can accommodate a variety of more dexterous equipment, including the Dual-Arm Work Module (developed by RedZone under a separate contract). This module provides two manipulators mounted on a five degree-of-freedom platform. The actuated motions are: a variable offset from center for each manipulator, an adjustable center pivot to orient the arms side by side or above and below each other, and a pivot at each manipulator allowing orientation from elbows-up to elbows down. An alternate tool which will be deployed from the heavy manipulator is a single manipulator. The lower weight of the single manipulator permits greater reach with the heavy manipulator to the sides and rear of the locomotor, but restricts dexterity to single arm tasks.

3.3.1 Locomotor

The locomotor frame is an aluminum weldment which supports wheel modules at each corner. Each wheel module has independent drive and steering motions providing an omni-directional capability. Mobility specifications of the Rosie locomotor are presented in Table 3.2.

Table 3.2 Rosie locomotor specifications

Width front axle retracted front axle extended	76 in 136 in
Height	42 in
Length	114 in
Obstacle Climb	4 in (max.)
Ground Clearance	6 in
Min. Turning Radius	about center
Max. Speed	2 ft./sec

The frame is a truss structure which supports wheel modules at each corner. Each wheel module has independent drive and steer motions providing an omni-directional capability. The wheels are controlled in any of three driving modes as illustrated in Figure 3.2.

- In *4-wheel steer* mode, the front and rear sets of wheels steer in opposite directions. Rosie drives along a circular arc (and its orientation changes with respect to the world as it drives). When the wheels are steered to either limit in this mode, Rosie rotates about the center of its locomotor.
- In *crab steer* mode, all four wheels steer in the same direction. Rosie drives along a straight line maintaining its orientation with respect to world.
- In *rotate-about-point* mode, the operator define a point in space – usually the tip of the boom – that Rosie is to rotate about. Rosie follows a circular arc whose center is at the defined point in space.




Figure 3.2 Rosie steering modes

The front two wheels are mounted on extensions which can extend the front wheel spread from 76 in to 136 in. The rear two wheels are mounted on a beam which pivots at the center to provide ± 2 in of vertical travel for obstacle negotiation.

Located within the locomotor frame is the hydraulic power supply, which is a 60 HP supply, providing 30 gpm flow at 3,000 psi for all robot motions. Filters and all locomotor valving are located in one of two side enclosures on the frame. The other side enclosure contains all control electronics for the system. At the rear of the machine is the tether reel which carries up to 200 feet of cable.

3.3.2 Heavy Manipulator

The heavy manipulator (or boom) is a four degree-of-freedom mechanism mounted near the rear of the locomotor deck. It can carry up to 1,700 lbs. with a 60,000 in-lb moment load throughout its work envelope, i.e., at a distance of 20 feet from the shoulder joint. Higher load capacities are available when the boom is at less than full extension. Load capacity is limited more by stability against tip over than by strength of the boom mechanism.

The heavy manipulator consists of four joints; a waist motion on the locomotor deck, a shoulder pitch, a forearm extension and a wrist pitch at the tip of the forearm. Each of the four joints has integral position feedback and is servo controlled based on operator commands. The joints can be controlled in either of two modes:

- *joint motion mode* allows the operator to individually control each joint on the heavy manipulator. This mode allows the operator to directly control the configuration of the system.
- *coordinated motion mode* allows the operator to control the Cartesian position of the endpoint of the heavy manipulator while the computer calculates the resulting joint positions. This motion mode is useful for tasks where the operator must move parallel to a floor or wall or is not moving in a confined area. Figure 3.3 shows a time lapse sequence of coordinated boom motions.




Figure 3.3 Rosie heavy manipulator in several configurations

3.3.3 Tooling and Auxiliary Services

A wide variety of tools or dexterous manipulators can be deployed from the heavy manipulator or locomotor deck. Highly accurate variable-speed motion control allows an operator to position tools quickly and perform work tasks effectively. Rosie's work envelope allows floor to ceiling reach with most tools. Both hydraulic and electric power are available at the boom tip to power tools. As much as 15 gpm at 3,000 psi and 20 amps of 120 VAC power are available. Virtually any

user-specified tooling can be deployed within these constraints, including those in the following table.

Table 3.3 Partial list of tools compatible with Rosie

Task	Tools
Process equipment & piping removal and size reduction	<ul style="list-style-type: none"> • manipulator(s) • pipe shear • abrasive saw/grinder • plasma torch
Structural steel and reinforced concrete demolition	<ul style="list-style-type: none"> • manipulator(s) • pipe shear • abrasive saw/grinder • plasma torch • jack-hammer/breaker
Concrete decontamination	<ul style="list-style-type: none"> • concrete scabblor • CO₂ pellet blaster • laser ablation system • coating sprayer
Debris (dry and wet) collection, packaging, and transport	<ul style="list-style-type: none"> • manipulator(s) • gripper • shovel • scoop • wet/dry vacuum

3.3.4 Control System

Control system functions are distributed across two primary computers (CPUs) — one in the console and one onboard the robot. The control console CPU displays status and sensor data coming from the robot, interprets signals from joysticks and other switches, and sends appropriate commands to the onboard CPU. The onboard CPU executes commands from the console by closing motion control loops, acquiring sensor data, coordinating axes, and activating video and other onboard equipment. Both CPUs perform continuous error checking and monitoring of communications between the robot and console. All software runs on VME-based 68040 Motorola CPUs under the VxWorks 5.1.1 real-time operating system. The two systems communicate via a dedicated ethernet connection in the tether.

All primary motions are servo-controlled enabling precise, variable speed motion control for positioning either by joystick command or by computer control. Feedback of absolute position of each steering actuator, wheel motor, and all four boom motions is provided by resolvers. The front wheel extensions incorporate limit switches so that the operator can easily discern whether they are extended; this information is also used by the control system in performing automated initialization sequences. The tether reel has limit switches so that the tether cannot be completely unwound from the reel, or wound on beyond the reel's capacity.

The status and health of the system are constantly monitored by various onboard sensors. These include hydraulic fluid temperature, pressure, and reservoir level. In addition, the voltage levels of onboard electronics components as well as their temperatures are also monitored so that a fault may be detected before it can cause a complete system failure. Control and sensing signals are monitored automatically and error checking is performed to ensure reliable communications.

3.3.5 Operator Console

A single operator stationed at the console (Figure 3.1) can control the Rosie worksystem. Primary system functions—locomotor, heavy manipulator, system power, tether, and cameras—are controlled using switches and joysticks on the desk top. Less frequently used functions and status information are accessed via a combination touchscreen/status display. No keyboard or mouse is required to run the system.



Figure 3.4 Rosie operator console

Three video monitors, with quad-splitting capabilities, display the onboard camera views. The operator can select any camera view for any of the monitors using the touch screen controls. In this way, each operator can configure the control console monitors to suit his or her particular preferences. Views can be changed at any time as needs arise.

The audio/video system takes multiple camera views and microphone inputs from the robot and displays them at the console. Rosie can support up to 10 cameras including the following:

- four cameras with remote focus, zoom, lights, and pan and tilt motions
- four cameras with remote lights and tilt motions (fixed focus)
- two cameras with remote lights (fixed focus)

All cameras are modular to allow easy replacement or relocation in order to accommodate different tooling or task requirements.

Rosie's operator is also provided with a complete set of feedback information (at the touchscreen/display) to support remote operations. Various other onboard sensors provide full system status and health monitoring. Additional user-specified sensors can be installed to provide remote monitoring of key environmental parameters.

3.3.6 Power and Telemetry

The power and telemetry subsystem allows power and signals to be transmitted from the console to the locomotor and routed onboard to the various sensors and actuators. A Power Distribution Unit (PDU) located between the console and robot provides a location to input site electrical power needed for onboard functions. A tether is used to transmit all power, control, and video signals to and from the robot. All signals from the console pass through the PDU and are combined with the power and routed into the tether. When operating in a contaminated location, the PDU can be located inside of containment, minimizing the number of conductor penetrations required through containment.

The heart of the electrical system onboard the locomotor is enclosed in a sealed box mounted on the left side of the frame. This enclosure houses transformers, control computing, power supplies, video modulation equipment, and heat exchanger units.

3.4 Cold testing of Rosie (Phase III)

3.4.1 Testing Program Summary

3.4.1.1 Objectives

During the development of Rosie, rigorous testing was identified as an prerequisite to successful deployment in a D&D project. Without exercising the system and understanding its performance characteristics, envisioning and planning in detail its use in a significant D&D project would be difficult. Without having comprehensively identified and rectified system weaknesses, few if any D&D planners could be convinced to consider its use. And finally, without significant operational history and data regarding overall reliability, acceptance of this equipment for use in contaminated or otherwise committing environments was impossible.

The primary objectives of the Rosie testing program were to:

- *Characterize the worksystem's capabilities.* This included Rosie's ability to maneuver to and within the work site, to move and position tools effectively to perform work tasks, and to adequate power and control tools.
- *Determine system deficiencies and identify design modifications or enhancements* that would improve its suitability to D&D applications. This included performing realistic tasks to ensure mission-relevant run time. All components, as well as overall system functionality, were scrutinized for reliability, ease of operation, and maintenance needs.

To accomplish these objectives we developed and executed a six-month testing program. Testing commenced after a two month integration and debugging period when Rosie was fully

operational. Based on experience during debugging, we commenced the testing program with a focus on several potential problem areas, as well as an overall concern for whether or not the system was meeting critical system requirements. In particular, the testing plan focused on:

1. Known or suspected problems:
 - unexpected/unintentional control system shutdowns
 - uncontrolled motion on any axis
 - unexpected loss of communication between console and robot
 - unstable performance of proportional valves
 - tether reeling system malfunctions
 - excessive noise on analog sensor signals and video
2. General areas of scrutiny, particularly with respect to reliability:
 - adequacy of hydraulic system cooling, particularly with maximum expected loads and maximum ambient temperatures
 - adequacy of onboard electronics cooling
 - structural weaknesses or excessive deflection
 - components with unreasonable short (less than 100 hours) operating life or mean-time-between-failure
3. General system characterization:
 - calibration/zeroing of all motions with respect to local (robot) references
 - speed ranges and general precision of each locomotor and boom motion
 - tracking accuracy of coordinated control modes on the heavy manipulator
 - operational difficulty relating to the operator console interface
 - adequacy of video/audio feedback for driving and tool deployment
 - appropriate and necessary regular maintenance items
 - safety hazards posed to personnel near the robot

Specific tests performed were selected to generally characterize performance and to stress all aspects of the system so as to reveal possible deficiencies. Tests were also designed to be representative of activities performed during a D&D project, though not necessarily in every detail. For example, task-appropriate tools were deployed, but mock-ups were used as surrogates for real facilities. However, the mock-ups were accurate with respect to material composition and overall thickness/dimensions.

Based on these criteria, general system characterization and three types of work tasks – concrete demolition, metal cutting, and material handling – were performed. Our overall philosophy during the work task tests was to, if possible, achieve a total system run-time on the order of what might occur in a D&D project. We also sought to handle the system as it would be in the field, so as to reveal the results of anticipated “abuse” from site conditions and operators. Under these testing conditions, we fully expected and intentionally sought system and component failures so as to identify any and all areas requiring modification or enhancement.

Late in Phase II, Rosie was transported to Oak Ridge National Lab's Robotics Technology Assessment Facility (RTAF) for system characterization and cold testing. The RTAF is a large high-bay with a ceiling over 40 feet high. The area was heated during cold weather, but not air conditioned during hot weather. Building temperature was reasonably well controlled, except when a nearby garage door was opened periodically for other operations in the building and this would locally heat/cool the system for short periods of time. The outdoor testing area was an asphalt driveway (leading up to the garage door) with standard curbs on both sides. The system was operated outdoors to perform general driving and all of the concrete demolition, and was used there regardless of temperature or weather conditions (including rain).

Most tests used a team of two operators, one to operate the system from the console and one to monitor the robot in the work area. The latter person made sure that no unauthorized personnel got near the robot while it was in operation and watched for any other unexpected problems. During all testing, system status was monitored, including hydraulic fluid and the electronics enclosure temperatures, hydraulic system pressure, input power stability, communications between the console and robot, and general command executions. Operators were provided with status information and warnings if parameters were out of defined standard operating range. Periodic (bi-monthly) visual inspections were also performed of the robot structure, hydraulic system components, and other onboard components to identify any potential failures.

3.4.1.2 General Results

At times, the Rosie system was operated without using tooling or performing a specific D&D task. These operations consisted of general driving, testing and verifying software changes, characterizing the motions and capabilities of the system, testing mechanical modifications, and demonstrating the system to visitors at the ORNL RTAF. An operations log was kept to document all operations of the system, such as what tests or modifications were performed, when they occurred, who was operating the system, and for how long. Run time recorded in the log included only those hours when the hydraulic power supply was running, regardless of whether or not the robot was actually moving or doing useful work. Overall, as of February 29, 1996, more than 180 hours of such operations were logged. This time does not include any time that the system was run and tested before it was shipped to ORNL; it is estimated that this totals an additional 40 to 80 hours.

The main tool deployed during the testing of the system at ORNL was the hydraulic concrete breaker. The next most frequently used tool was the Schilling Titan II dexterous manipulator, which also deployed smaller tools. The breaker was used a total of 80 hours on 23 different days with the longest continuous run being 4.5 hours. The Schilling arm was used alone for a total of 8.2 hours over 3 days. The electric grinder (deployed from the Schilling arm) was used for a total of 22.5 hours over 4 days. Overall the system has been run on 49 different days between September 1995 and February 1996. The average run time per day has been 3.7 hours and the largest amount of run time accumulated on one day was about 8.5 hours.

While at ORNL, Rosie was operated under many different environmental conditions: indoors and outdoors in both winter and summer. Ambient temperatures ranged from about 30°F to 95°F. The outdoor conditions ranged from hot, sunny days to cold, overcast days with constant rain. The system was also stored outdoors for several days at a time. None of these conditions caused any noticeable problems, other than small amounts of surface rust on some of the unpainted ferrous

parts. The locomotor has been driven on concrete, asphalt, and even grass with no apparent problems, though on very hot days the wheels leave indentations on in asphalt.

3.4.2 General System Characterization

A series of tests was performed to quantify Rosie's motion characteristics including speed ranges and smoothness of motion on each axis of the locomotor and heavy manipulator. Some tests were also performed to characterize the position accuracy and repeatability of the heavy manipulator motions under cartesian control

The locomotor is omni-directional: it can follow curvilinear paths of any curvature and can change direction instantaneously. Its driving speed ranges from about 0.2 feet per second to a maximum speed of 2 feet per second. Driving speed is proportional to the deflection of a console joystick from its nominal or "home" position; at a constant driving speed, the distance driven depends on the length of time the joystick is held in that position. Tests were performed to determine the smallest incremental distance the locomotor could be commanded to move using joystick inputs in both four wheel steer and crab steer modes. The smallest distance that the locomotor could be moved in either steering mode ranged from 1 inch to 1.75 inches. No attempt was made to determine the limiting factors controlling the size of this increment. It may be limited by the components in the hydraulic system or it may be related to the tuning of the servo control loops which are controlled by the system software. If the limits are in the servo loops, further tuning may improve the performance.

The curbs on the sides of the driveway were used to test the ability of the locomotor to drive over obstacles. The curbs are about five inches tall and almost vertical. The locomotor had no difficulty in driving over these curbs, either straight on, using two wheels at once, or at an angle, using just one wheel at a time. In addition to the curbs, many smaller pieces of broken concrete debris were driven over during the testing. No difficulty was encountered in doing this, and no problems resulted.

The heavy manipulator axes have maximum speeds of 3 degree per second on the three rotational axes and 6 inches per second on the extension axis. The minimum speeds for these axes are very close to zero but were not measured exactly. This is because joint speeds are non-linearly mapped to the boom control joystick motions; there is a step function drop from very slow to zero near the joystick center position.

The waist rotation axis had exhibited some undesirable characteristics in the initial system trials, so it was examined more closely than the other axes. At slower speeds this axis did not seem to be able to produce smooth motions, but instead exhibited a cogging type motion. This characteristic interferes with operations when the boom-waist motion must be used to precisely align the boom tip with an object, or if a constant slow-speed waist motion is required. We suspect that this is due to the cogging characteristics of the motor that drives the axis and other frictional forces in the mechanical components (bearings, drive train, mechanical stiffness, etc.) that comprise this axis.

To characterize the motion of this axis, the output of its position sensor (resolver) was measured versus time while the axis was commanded to rotate slowly. This showed that a speed variation occurred a little more than once per degree of rotation. Next the servo system was disabled and the axis control valve was driven open loop. This produced virtually identical results. The drive motor

for the axis was then disconnected hydraulically (so that there would be no hydraulic influence) and the axis was moved manually. The cogging was readily apparent to the person moving the axis. It manifested itself as a frictional variation as the axis is rotated and was again, a little more than once per degree. Finally, the preloading on the pinion gear which drives the axis was reduced to zero and the axis was moved manually again. This time there was no cogging and the axis moved freely and smoothly in either direction.

These results indicate that these variations are produced by uneven frictional characteristics of the main drive gear train for this axis. The frictional variation is high enough that the servo system on this axis cannot compensate for it. At higher speeds the inertia of the axis tends to smooth out the variation and it is much less noticeable. At low speeds (less than one degree per second) the inertia cannot smooth out the frictional variations and the result is a speed fluctuation every 0.96 degrees, which is 4 fluctuations for each tooth on the main drive ring gear. In addition to affecting the smoothness of the waist rotation velocity, this variation affects the position resolution of the axis.

To characterize the linear motions of the heavy manipulator under Cartesian control, a plumb bob was attached to the end of the boom and held just off the floor. Using the joystick, the boom tip was commanded to move horizontally in a straight line for about 10 feet. The start and finish points were marked and the motion was repeated up to five times. The plumb bob consistently returned to the start and finish points within 0.5 inches. The linearity of the motion was not measured explicitly, but a spot check in the middle of the 10 foot range indicated that it was as accurate as the end points. These motions were measured with the boom extended out toward the front of the locomotor and also with it extended out to the left side of the locomotor. The results were the same in both cases.

The resolution and sensitivity of the motions of the heavy manipulator were measured by repeatedly positioning a plumb bob (hanging from the wrist joint) over a marked point on the floor. To start, the operator tried to position the plumb bob over the mark by teleoperating the heavy manipulator only. The plumb bob was then moved away from the mark to an arbitrary position, and another attempt was made to position over the mark. This test was performed at least five times each using both joint and coordinated modes with two different operators. The results were consistent; each operator could position the plumb bob within a 1 inch diameter circle at the mark using either mode of control.

Although payload capacity testing was not performed at ORNL, some testing did occur during fabrication. Before delivery to RedZone, the heavy manipulator underwent an acceptance test at PaR' System's (the manufacturer) facilities. The manipulator was rigidly mounted to the floor for the tests, so that tip over was not an issue. The system was shown to be capable of lifting a 2,000 lb payload with either a 60,000 in-lb forward moment or a combined 25,000 in-lb forward moment and 43,000 in-LB side moment load. After the heavy manipulator was installed onto the locomotor, it was tested at RedZone's facilities using a 1,500 lb payload which had a 40,000 in-lb forward moment. The manipulator was also tested with a 3,000 lb load, which was lifted only with the boom in the retracted position. The heavy manipulator had no trouble lifting any of these test loads. Further, these tests showed that the design estimates for tip over were slightly on the conservative side.

3.4.3 Cold Testing of D&D Tasks

3.4.3.1 Concrete Demolition

The majority of cold D&D testing consisted of demolishing concrete structures using a hydraulic breaker (jackhammer). The breaker was mounted on the tooling plate at the tip of the heavy manipulator such that its tip was approximately four feet from the payload mounting surface. The one-way solenoid valve that controls the breaker was also mounted on the tooling plate and plumbed to the manifold at the end of the boom fly section with hoses. The valve was controlled from the operator console using a foot switch to turn the breaker on or off.

The breaker could be positioned anywhere from in-line with the boom to perpendicular to the boom (i.e., pointing down toward the ground). Two "wing" cameras, each with a tilt motion, were also mounted on the tooling plate, one on each side. These two cameras were separated by about six feet and angled toward the tip of the breaker to provide some depth perception when positioning the breaker tip on the target concrete.

Several "Jersey" highway barriers and precast culvert sections, all made of steel-reinforced concrete, were used in the demolition tasks. For safety, all demolition was performed outside the RTAF. The typical procedure was to position the locomotor near the object to be demolished and then use the heavy manipulator motions to place the breaker in the desired position beside or above the barrier. The breaker was then pushed against the target surface to ensure that adequate pressure was maintained during demolition. There was no force feedback in the system, so the operator would generally determine that adequate force was being exerted by observing motion of the target object, the locomotor, or both (this procedure is commonly used with construction equipment in conventional demolition.)

Six Jersey barriers (approximately 10 ft. long, 3 ft. high, and weighing 2 tons each) were demolished using this procedure described. The barriers were reduced to pieces approximately 6 inches cubed or smaller over approximately 80 hours of operation spread over 23 days. Reduction of a single barrier into 1 cubic foot pieces could be accomplished in about 4 hours; further size reduction and separation of these pieces from the reinforcement bars took the remainder of the time. It was felt that this latter operation could be more efficiently accomplished using alternate tools, such as a pulverizer, or by fitting the breaker with a broad, crushing type tip.

Teleoperated control proved to be an acceptable method of operation for concrete demolition. After a short period to learn how to judge distances using the two wing cameras, an operator could readily position the breaker in a desired location. For example, it was not very difficult to place the tip of the breaker in the middle of the top edge, which is approximately 4 inches wide, of one of the Jersey barriers. Video feedback was found to be critical to successful and efficient breaking; the views provided by the boom tip cameras were adequate, but a close-up/zoom view would allow more precise alignment of the breaker tip and generally improve efficiency. Audio feedback from the two microphones located on the robot also proved to be very useful in concrete breaking. It let operators know when they were solidly breaking concrete and when they needed to apply more pressure to the breaker tip, or when the tip had broken through the concrete.

One problem encountered during concrete demolition was lack of awareness that the locomotor had tilted in reaction to forcing the breaker into the work piece with the boom. The shoulder joint of the heavy manipulator generates enough force to lift the front wheels of the locomotor well

over one foot of the ground, even when the boom is fully extended. Although this would not likely lead to tip-over, it can result in unexpected and sudden motion if the object being pushed against suddenly gives way. None of the onboard cameras give a very obvious indication of this tilting motion; addition of a tilt sensor on the locomotor and a display at the console may help give the operator a sense of when the locomotor is tilting.

The hydraulic breaker used during the concrete demolition required frequent lubrication (every 2 to 5 hours), which had to be done manually. In a future upgrade, some type of automatic lubricating system should be incorporated into the breaker to allow completely remote operations.

3.4.3.2 Metal Cutting

Rosie was used to remotely cut metal tubing and angle stock. To perform this work, a Schilling Titan II dexterous manipulator was mounted on the heavy manipulator tooling plate. Hydraulic supply and return lines were run from the manifold at the boom tip of the boom and routed across the wrist axis to the manipulator. The Titan II's electronic control box was also mounted on the tooling plate and connected to Rosie's electronics enclosure via a cable that is permanently installed in the boom's cable track. Conductors in Rosie's tether completed the electrical connection between the manipulator local electronics and its master controller at the operator console.

A welded framework of stainless steel angle stock (1" x 1"; 1/4" thickness) was used as the target object for these tests. The manipulator picked up an electric grinder that was modified with a T-handle that could be held in the manipulator's parallel jaw gripper. Two grinders, one with a 4" blade and one with a 9" blade were used. Cutting operations were performed in two different ways. In one method, the locomotor and heavy manipulator first positioned the Titan arm near the work piece after which the Titan made all subsequent motions. In the other method, the swing motion of the heavy manipulator was used to move the Titan into alignment with the pieces to be cut first. After each piece was cut, the heavy manipulator repositioned the Titan to cut the next piece. This method was especially useful when a row of parts, such as several vertical pipes, needed to be cut.

About 22.5 hours of cutting was performed over a four day period on the stainless steel test piece. Some additional cutting (approximately 3 hours) was performed on reinforcing bar steel exposed during concrete demolition. Overall, nearly half of the time was consumed picking up the tools with the manipulator, changing blades, and moving the mock-up around. Only three hours of cutting was performed using the small grinder which exhausted two blades producing four crosscuts on the angles in the framework. With the large grinder, five blades were consumed or broken while making about 24 crosscuts in the angles and three plunge cuts into the main frame supporting the angles.

The greatest difficulty encountered in cutting metal parts with the grinder was the inability to see exactly where the cut was being made and what the angle of the tool was relative to the work piece. This was due to the placement of the cameras at the boom tip. No camera was mounted directly on the manipulator for this testing; a camera on the Titan wrist would certainly have improved the situation, but the operators still felt that it would not solve the problem completely.

The tool itself gets in the way of the camera views if the cameras are too close to it. A camera farther away and to the side of the manipulator with a zoom lens may be a better solution.

The grinders themselves were difficult to pick up when lying on the ground because of the clearance around the T-handle and the guards on the grinders. When human observers are present, guards are needed for safety, however, for truly decommissioning tasks, removing the guards would allow a better view for picking up and orienting the tool and would allow faster blade changes.

During testing, stainless steel abrasive grinding wheels were not available, and therefore cutting operations were not highly effective. The rate of blade consumption/breakage experienced during this test is considered exceptionally high; this was partially due to use of suboptimal blade materials (i.e., non-stainless type blades wore down very quickly). Lack of force feedback, which can help an operator prevent binding of the cutting disk in the work, and inadequacy of camera views also contributed to excessive blade wear and breakage.

Rosie's heavy manipulator and locomotor provided an adequately stable, rigid base platform for the Titan II during its operations. Overall, Rosie's ability to support this type of tool deployment and the effectiveness of remote cutting with the abrasive grinder were quite acceptable, although longer duration testing with appropriate blades should be performed if possible.

3.4.3.3 Debris Removal

While the Titan manipulator was mounted on the boom, some brief testing was performed to determine how useful it would be in picking up debris left over from the concrete demolition. The Titan II gripper can only open slightly more than four inches which restricts it to picking up only pieces that are roughly that size. Larger pieces (up to around 8 or 9 inches) could be picked up if those pieces had a smaller protrusion which would fit within the gripper jaw capacity. Positioning the arm to get a good grip during this type of work was very tedious, and it was easy to squeeze too hard with the gripper and break the piece being grasped.

3.5 Modifications and enhancements of Rosie

Several failures and malfunctions occurred during operation and testing of Rosie-O at ORNL, none of which were catastrophic or insurmountable. Many were corrected by modification of Rosie-O; others were not implemented on Rosie-O but instead accounted for in the Rosie-C design. Motivations for and descriptions of these changes are described below.

3.5.1 Hydraulics

During the course of operations, several proportional valves became uncontrollable. They were replaced with valves from another vendor which have since worked adequately, though performance is still less than desired. The Rosie-C design uses servo valves instead of proportional valves to alleviate these problems.

The original proportional valves used in the system also allowed the heavy manipulator to droop when power was removed from the system. After exercising the system with this feature for some time, it was determined that keeping this axis locked in position would be a safer alternative when the machine is shut down. The replacement proportional valves were therefore specified to

implement this locked shutdown mode. The Rosie-O boom now stays in position when power is removed; very slow droop (on the order of many hours or even days) still occurs due to leakage in the hydraulic cylinders that power the shoulder axis. In the Rosie-C design, counterbalance valves have been added to the shoulder cylinder circuit to lock the boom in place, even in the event of a hose failure.

The cooling system for the hydraulic fluid seems adequate under normal conditions, but could be insufficient when the ambient temperature is high and the system duty cycle is low (i.e., it is spending a lot of time idling). The Rosie-O cooling system only cools fluid that returns from the actuators, so when the system is idling, very little fluid flows through the heat exchanger. Sufficient cooling margin is of particular concern because Rosie's hydraulic fluid (Houghtosafe) is a water/glycol mixture that loses much of its lubricity at temperatures above 180°F. Though no failures were encountered, Rosie-O's cooling system is marginal. Rosie-C has a separate ("kidney") cooling loop that continuously cools *all* fluid regardless of hydraulic power demand. This should improve the cooling performance to be sufficient even under extreme operating conditions.

Several other changes were made to the hydraulic system:

- *Redesign of the hydraulic reservoir.* With a larger reservoir, the amount of time that entrained air bubbles can escape from the fluid is increased, as is the efficiency of heat transfer out of the reservoir. The new design also improves strength, ventilation within the tank, connections to it and accessibility for maintenance.
- *Changes to improve the life of the main pump and other hydraulic components.* These include faster flow compensation (which also squelches pressure spikes that reduce seal integrity), improved fluid filtration, simpler piping connections to the pump, and better pump shaft seals.
- *Improved design of hydraulic swivels on the wheel modules* for simpler, more reliable connections.
- *Improvements to cooling and filtration system.* These include the addition of a kidney loop that provides better heat dissipation, easier access to fluid filters, addition of a recirculation pump and improved air flow across the heat exchanger.

3.5.2 Locomotor

The rear rocker arm axle has two hydraulic rams mounted inside it for the purpose of locking the axle in position, thereby reducing the compliance of the locomotor as a base platform for manipulation. The hydraulic circuit that locks these rams consists of two valves, one to supply and release pressure, and one to isolate the two rams from each other. The isolation valve was found to work correctly in one direction but leaks slowly in the other causing the locomotor to tilt in one direction when the rocker arm is in the locked position. A new hydraulic circuit was devised to circumvent this phenomenon and is being implemented on Rosie-C.

Though the locomotor is capable of driving over a 4-inch high object, it can run over a higher obstacle if the front edge is angled or sloped. Occurrences such as this resulted in minor damage to the underside of the electronics enclosure. In the Rosie-C design, the bottom of the enclosure has been reinforced to resist such damage.

Small stress cracks have developed in all of the corners of the hydraulic enclosure where the aluminum parts were welded together. These cracks were only in the outer skin areas (near where the doors attach) and thus were not considered critical. On Rosie-C these areas have been reinforced.

A crack was discovered in a weld in the right front section of the locomotor frame where one C-channel attaches to another. Since this area is not highly loaded and should experience compressive loads only, the cause of this crack is unclear, though its presence is not a safety concern. Regardless, triangular gussets have been added to this location in the Rosie-C design.

Each side of the extensible front axle uses an E-chain cable carrier to guide hydraulic lines and electrical cables between the locomotor frame and the wheel modules. These E-chains are mounted above the cylinders that move the extensions in and out. When a joint in one E-chain segment broke, the E-chain drooped lower than normal, came into contact with the cylinder, and was eventually damaged by cylinder motion. In the Rosie-C design, the E-chains are mounted below the cylinders to avoid such problems.

The initial integration and testing of the Rosie system was performed at RedZone's facility in Pittsburgh during the summer months of 1995. When ambient temperatures were above 90°F, the electronics enclosure internal temperature rose higher than expected. Since this enclosure is totally sealed, air-to-air heat exchangers are used to remove heat from it. The air inlets for these heat exchangers are located on the bottom of the enclosure. Investigation revealed that hot air coming from the hydraulics radiator was flowing directly into the electronics enclosure heat exchangers, reducing the temperature differential they were experiencing, hence their ability to transfer heat out of the enclosure. With lower ambient temperatures, this is not a problem, but in a hot summertime ambient, it is not acceptable. An interim solution, which has worked adequately, was addition of a temporary air deflector to re-direct hydraulics radiator exhaust away from the electronics enclosure. The solution for Rosie-C is to reverse the airflow through the hydraulics cooling radiator away from the electronics enclosure. In addition, the heat load inside Rosie-C's electronics enclosure is significantly lower.

The following additional minor changes were made to the locomotor:

- *Wheel module re-designs.* The resolvers which measure wheel module steering angle were repositioned, terminal blocks were substituted for the electrical connectors and covers added to the brackets. These changes provide better clearance, hence more protection. In addition, the gearing ratio in the belts that link the resolver shaft to the steering axle was changed to prevent electrical wrap-around. The new hydraulic swivels on wheel motor assemblies now accommodate pass-through of the wheel motor resolver cables.
- *Structural improvements.* Mounting of several components, including the electronics enclosure and the hydraulic reservoir, was modified and access holes added to make maintenance and assembly easier. Gussets were added in locations where cracks in the Rosie-O frame appeared.

3.5.3 Heavy manipulator

Very early in the concrete demolition testing, the wrist axis load mounting platform failed mechanically. Concrete breaking was not anticipated when this platform was designed and the

additional loads imposed by this tool caused the failure. A new, stronger load mounting platform was designed, built, and has been in use on Rosie-O for over 150 run time hours.

Cracks appeared in the welds at the front corners of the base of the heavy manipulator. These were likely due to the extra loads imposed by concrete breaking, as this structure was not originally designed for such loads. Steel reinforcements were fabricated and bolted in place to relieve the stresses in the welded areas by transferring loads directly to the locomotor frame. Rosie-C's heavy manipulator base is constructed of steel instead of aluminum to avoid this problem.

During a concrete breaking test, the front axle of the locomotor was raised about one foot off the ground and suddenly fell when the breaker tip slipped off the piece being jackhammered. The resulting dynamic loads experienced by the counterweight support brackets caused their welds to fail catastrophically and all four counterweights fell to the ground. Subsequent examination of the failed pieces indicated that the welds were defective and did not meet specifications. They were repaired by ORNL personnel using full penetration welds and the counterweights were reinstalled. No further problems were encountered. In the Rosie-C design two (as opposed to Rosie-O's four) counterweights are secured to the boom with fasteners instead of being hung in place.

The attachment bracket for the rod end of the middle boom section cylinder experienced some deformation during testing before the system was shipped to ORNL. The attachment was designed to support a 2,000 lb. load at the boom tip. During testing at ORNL, the boom tip was pushed against a wall, and the rod attachment was loaded with the full output of the cylinder, which is over 9,000 lb. The attachment bent plastically. A reinforcing bracket was designed, fabricated, and welded onto the attachment structure by ORNL personnel. In addition, a pressure regulator was installed in the boom extension hydraulic circuit. This regulator limits the pressure available to the cylinder to 2,000 psi, and reduces the maximum force to just over 6,000 lb. In the Rosie-C design, the attachment point has been structurally reinforced so that it can handle the maximum output force of the cylinder, obviating the need for a pressure regulator.

The heavy manipulator waist rotation axis has some undesirable friction characteristics that affect the smoothness of motion at slow speeds. This friction is generated in the main drive gears and is high enough that the servo loop around this motion cannot compensate for it. Though some alternative designs have been considered (lapping the gears for a better fit, stiffening pinion gear mount, reducing the width of the pinion gear, using a belt drive instead), none have been implemented on either Rosie-O or Rosie-C. A solution to this non-linear behavior is still needed, particularly if Rosie is to be operated as a robot, i.e., in automatic mode under the control of a computer.

Several other minor changes were made to the heavy manipulator design.

- *A lifting hook* was added to the underside of the boom just before the wrist actuator.
- *Functional changes.* The shoulder cylinder pressure transducers were eliminated since the method for calculating boom tip load based on measured cylinder pressure proved to be unreliable. Removable mechanical stops were added to the shoulder joints to prevent the boom from hitting the locomotor deck.

- *Changes to facilitate maintenance.* External grease fittings were added to the waist rotation bearing. The boom-top E-chain was changed from closed type to open type for easier decontamination of the hoses and cables within it.

3.5.4 Tether Management System

The original capstan mechanism to keep tension on the tether as it is reeled on and off the tether reel did not work properly. It was electrically actuated and used two "V" profile belts driven with pulleys and a gear train to tension the tether. This mechanism had too much internal friction and required constant adjustment to work properly. Breakage of the V-belts was also a recurring problem. A completely new design was conceived, fabricated, and installed on Rosie-O. Over 30 hours of operation have been logged and the new design is performing properly.

The tether management system was substantially re-designed; the new design was implemented and tested on Rosie-O. The capstan drive was completely re-designed. The single electric motor drive was replaced by two hydraulic motors, each driving a V-grove pinch roller on one side of the fair lead. This change improves the coordination of the tether drum and the capstan and allows better regulation of tether tension, as well as eliminating the complex gear train that transferred power from one side of the fair lead to the other. Lead-in rollers were added to help guide the tether through the capstan.

3.5.5 Control System

To enhance safety, a pendant style E-stop switch was added to the locomotor. A person standing near the locomotor can use this switch, which is on a 20-foot cable, to shut down the system in the event of an emergency.

The front wheel extensions currently do not have any position feedback; once they are actuated they are presumed to reach a limit of travel. The control system therefore only infers their position. Limit switches have been added to the Rosie-C extension cylinders to provide positive feedback. With these devices, the control system will know if the extensions are all the way in, all the way out, or at some intermediate position.

During initial testing at RedZone the system randomly shut down several times with a software-generated E-stop. Several causes of this problem were identified and corrected. It was also discovered that if the onboard controller is left running and the console software is rebooted, the controller crashes. This problem was easily fixed once it was found and has not recurred.

When the shoulder axis moves, the wrist axis is supposed to "track" its motion such that the tooling plate attitude relative to the floor remains constant. In the original implementation of the software, this function used velocities as a basis and was always running. This has been changed so that the wrist more closely follows the shoulder motion and the tracking can be turned off or on, as the operator desires.

A teach/playback mode was implemented for the heavy manipulator. It allows the operator to record the sequence and timing of a series of motions (as they are executed) for subsequent playback in either absolute or a relative mode. These features are accessed via the console touchscreen.

3.5.6 Console

Prolonged use of the locomotor and manipulator joysticks caused operator hand fatigue, which appeared to have two causes. The joysticks were not ergonomically positioned on the desktop, making for an uncomfortable working position. Further, the joystick centering springs were very stiff, thus a substantial amount of effort was required to move them. In the Rosie-C console, the shape and layout of the console has been completely remade and the joysticks use lighter centering springs.

The deadband at the center of the joystick travel has not yet been optimized. As a joystick is moved from its furthest deflection toward the center ("home") position, its output varies linearly from a maximum value down to some minimum, then falls sharply to zero before the joystick reaches home. This was changed such that zero output is reached and held within a deadband near the center of the joystick's range of motion.

Several operators recommended changes in the layout and functionality of controls and indicators on the control console touchscreen/display. The functions added include teach/playback controls, an axis jog function, a axis position display, camera assignment presets, and enhanced system status displays.

3.6 Artisan Task Space Scene Analysis System

Cold testing in Phase I indicated that while mobile worksystems, such as RWV and Rosie, offer a means to perform D&D tasks remotely and with substantially less risk of human exposure to radiation, dangerous chemicals or occupational hazards. However, there is great inefficiency in getting work done when the machine operator is removed from the work site and linked to it with closed circuit television. The productivity loss is well known in the remote operations community and despite genuine motivation for teleoperated control of equipment, operator interfaces to remote hardware have improved surprising little over several decades of research, development and practical application. In some cases, overall productivity of a remote worksystem is only 10% of analogous "hands-on" or operator on-site modes.

Artisan is a semi-automatic perception system that combines range sensors, object modeling and analysis software, and an operator interface to create a 3-D model of a robot's work area – *in situ*, i.e., as it is discovered by the robot. This paradigm, known as task space scene analysis, provides a much richer understanding of complex, interior work environments than what can be gleaned from conventional 2-D camera images. Scene analysis techniques that map space as occupied and unoccupied regions of space are sufficient to plan collision-free motion trajectories. Artisan goes beyond this, however, to provide additional information that enables automatic task execution. Modeling of the world has advantages over 3-D viewing and mapping systems because it allows both the robot and its operator to understand the not just the layout, but also the context of the work space and the objects in it. Since objects in the scene are *recognized*, semantic information such as function, likely material type and even manufacturer, is acquired along with shape information. This creates opportunities for pre-planned robot task execution specific to particular object types, a capability not offered by systems that generate 3-D maps of occupied space only. Artisan thus creates the opportunity to elevate operator interaction with D&D worksystems from man-in-the-loop teleoperation to supervisory/telebotonic control and bring the benefits of safeguarding, finesse, repeatability and productivity of automation to bear on tasks that have been traditionally relegated to remote control.

The principal objective of task space scene analysis (TSSA) is construct a 3-D model of the work zone around a robot or other remote worksystem so that automatic control techniques, such as trajectory planning, collision avoidance, and scripted motion sequences, can be used to speed task execution relative to conventional teleoperation. Prior and related ongoing TSSA efforts within RTDP emphasize human interactive approaches in which the operator matches and adjusts the registration of object models to observed data using graphical overlays. To complement these, we developed a semi-automatic TSSA system, which we call Artisan. (Early in Phase II, a survey of related work was conducted and general requirements formulated Results of these efforts are presented in the Appendices of our Phase II Topical Report - see Executive Summary for DOE document number.)

An Artisan session begins when the operator instructs the system to acquire several images of the scene. Artisan uses a scanning laser range finder which produces both a range image and a reflectance image, both of which are displayed on the workstation screen. Since the field of view of the sensor is usually larger than the area that the operator wishes to work on, he restricts the system's attention to a particular region of interest by drawing a box around it. The operator then indicates what objects Artisan should expect to find in that region of interest by selecting from a menu of pre-defined object types and sizes.

Two different object recognition algorithms have been developed for Artisan. The first method (Quadric/Planar Segmentation and Matching, or "QPSM") segments the 3-D surfaces into planar and quadric patches and matches the resulting scene description to analogous descriptions of object models in a database (developed off-line from CAD descriptions of objects). The other method (Free-form Object Recognition Method or "FORM") is based on a technique known as geometric indexing. In this case each point in a 3-D collection of points is transformed into a 2-D representation, called a spin image, that describes the spatial relationship of all the other points to the selected point. Each point in a 3-D data set – whether from a range sensor or from a CAD model – has a corresponding spin image. The first pass in the recognition algorithm is comparison of random samples of spin images representing the scene data to spin images of models in the object database, which results in a few plausible correspondences between data and library model. In the second step, plausibility is further tested and the best matches are refined using a modified iterative closest point (ICP) algorithm that outputs an optimal estimate of the recognized object's location and orientation in the task space. QPSM was developed in Phase II and FORM was developed in Phase III.

Artisan exports its results into an interactive robot simulation package, TeleGRIP™ by Deneb Robotics, Inc. TeleGRIP keeps track of the geometry of the entire task space including all recognized objects, all *a priori* information and the remote worksystem itself. TeleGRIP provides visualization capabilities that allow the operator to view the task space and because the model is synthetic, the operator can change the perspective to see behind and underneath objects even if it is impossible to move a real camera into that position. More importantly, the operator can command the robot to perform a task by running a simulation. TeleGRIP provides several intrinsic motion planning capabilities as well as hooks that allow additional task programming scripts (developed off-line) to be executed at run time. For example, in a decontamination operation, the operator would like the robot to follow the surface of a pipe while maintaining a particular stand off distance. The operator selects a starting point for the end effector and the appropriate trajectory is calculated automatically. During simulation of task execution, TeleGRIP will warn the operator of any collisions or motions that exceed joint limits of the worksystem. Once satisfactory performance is obtained, the operator can "download" the resulting tool paths to the real worksystem.

3.6.1 Fundamental Aspects of Artisan

The Artisan software system consists of seven modules, the functions of which are summarized in Table 3.4.

Table 3.4 Artisan Modules

Artisan Main Menu (AMM)	Allows the operator to acquire images, to choose object models from a library that should be found in a selected region of interest, and to verify object matches
Range Image Module (RIM)	Acquires range/reflectance images; applies temporal and spatial filters to range images; generates meshes from range data
Range Image Display Module (RIDM)	Displays range and reflectance images; allows selection of a region of interest by the user; provides image enhancements for viewing purposes
Mesh Display Module (MDM)	Displays meshes and models generated by other modules; allows the user to resize matched object models and view objects from arbitrary viewpoints
Object Recognition Module (ORM) <i>used in QPSM only</i>	Segments mesh data into planar & quadric patches; matches the operator-selected models to mesh patches; resizes models to match mesh data

Table 3.4 Artisan Modules

Free-form Object Recognition Module (FORM) <i>used in FORM only</i>	Geometrically indexes mesh data into spin image representations; matches spin images of operator-selected models to spin images generated from data; refines pose estimates using a modified iterative closest point algorithm
Presentation & Robot Control Module (PRCM)	Receives object models after the operator has accepted them; performs path planning based on the models and interfaces to a real robot (built around TeleGRIP)

All Artisan software modules with the exception of range image data acquisition and filtering run on an Indigo² Extreme workstation from Silicon Graphics, Inc. (Though we are versed in RTDP's GENISAS communications software – several of our software engineers completed a tutorial presented by representatives of Sandia National Lab – we opted to use CMU's *TCX* software for interprocess communication since it is more flexible in its handling of multiple, variable and user-defined data structures, as well as being somewhat more efficient.)

When Artisan is running, the workstation screen appears like the layout shown in Figure 3.5. An *Artisan* session is controlled through a main menu (upper right-hand corner of Figure 3.5) that contains buttons for access to other Artisan modules.




Figure 3.5 Overall Artisan screen layout

The sensors used for Artisan development included two 3-D scanning laser rangefinder manufactured by Perceptron, Inc. The older device has a fixed field of view (60 degrees in both

elevation and azimuth) and acquires 256 x 256 pixel range and intensity images at a frame rate of 2 Hz. This model has an operating range of 2 to 40 meters and range precision of 5-7 cm. The newer LASAR 5000 (Figure 3.5) has a slightly larger maximum field of view, but the actual image field of view is programmable and the image size can be up to 1028 pixels in azimuth by up to 2048 pixels in elevation. Both generate amplitude-modulated laser signals using near-infrared (810 nm and 835 nm for the old and new models, respectively) laser diodes. The laser beams are scanned through the sensor field of view by nodding mirror and a rotating polygonal mirror, reflected of a target, gathered by the receiver optics and focused onto a detector. Phase difference between the detected signal and the transmitted beam is proportional to distance traveled to and from the target and is used to generate a range image while the amplitude of the return signal is used to generate a reflectance image.




Figure 3.6 Scanning laser rangefinder

The *Artisan* software makes no assumptions about the specific sensor that gathers the 3-D surface data. For example, different 3-D imaging sensors may have differing geometries (field of view, depth of field) and differing statistical characteristics (resolution, repeatability). All interaction with the 3-D imaging sensors occurs through a well-defined interface that hides low-level details about the sensor from the remainder of the *Artisan* system. With this 3-D imaging system abstraction, new sensors can be easily integrated into the system. The system has been tested with an older model Perceptron laser scanner, a newer Perceptron LASAR 5000, a K2T structured light rangefinder, and data files generated by the Coleman Coherent Laser Radar (CLR) sensor.

Figure 3.7 shows a pair of range and reflectance images (on the right and left sides of the figure, respectively) taken of one of our mock-ups. The scene consists of a series of cylindrical pipes of various dimensions, conduits, process vessels and an electrical panel mounted in a framework of I-beams.




Figure 3.7 Range image (left) and reflectance image (right) of experimental testbed

To focus the system on specific objects in the scene to be analyzed, the operator selects a rectangular region of interest in the image. This limits the amount of data to that which is important to recognizing the specific object. The region of interest can be selected in either the range images or the reflectance image dragging the mouse to form a rectangle, as shown in Figure 3.7. Once the region of interest is determined all of the range image points within it are filtered by a temporal averaging of an image sequence to reduce noise caused by temperature drift of the sensor. This is followed by the application of a spatial smoothing filter to eliminate outliers while preserving range discontinuities. The resulting range image is then converted from spherical sensor coordinates to Cartesian world coordinates (x,y,z) .

3.6.2 Quadric/Planar Segmentation and Matching Method

Quadric/Planar Segmentation and Matching (QPSM) was developed as an initial approach to semi-automatic TSSA in Phase II. Our work in Phase III created a newer, more general method, obviating the need for QPSM. A very brief explanation of this method is presented in the following paragraphs; a more detailed explanation of the underlying principles and algorithms of QPSM can be found in Chapter 4 of our Phase II Topical Report or in [14].

When using the QPSM method, Artisan creates a Cartesian mesh from the range data in the region of interest then segments it into planar and quadric patches. Planar patches indicate flat surfaces, such as walls, floors and I-beams, while quadric patches are parts of curved surfaces like pipes, tanks and valves. Figure 3.8 shows an unsegmented mesh formed from the range data of Figure 3.7 and the mesh patches resulting from segmentation. Artisan automatically tries to match these surface patches to geometric descriptions of the objects that the operator identified earlier and displays the results. For each object recognized, the operator can either accept or reject what Artisan recognized. Accepted objects then appear in an IGRIP "world" model at the locations that Artisan has calculated. The entire process from range image acquisition through object recognition takes approximately three minutes on the SGI Indigo² workstation.




Figure 3.8 Cartesian mesh formed from range data

These capabilities were first successfully demonstrated in early March 1995. The demonstration used a mock-up containing objects representative of selective equipment removal scenarios (pipes, tanks, conduits, I-beams, etc.) as the target scene. It also included a simplistic graphical programming example to illustrate the viability of the approach for telerobotic selective equipment removal. In this case, an industrial manipulator was situated in front of the mock-up such that it could reach most of the objects in it. The manipulator was also simulated in IGRIP. Objects recognized by Artisan then appeared in IGRIP at the correct location relative to the simulated manipulator, which was then commanded to follow a trajectory along one of the object's surfaces (to simulate a spraying or washdown decontamination process). The simulated manipulator performed the task successfully, i.e., there were no collisions with other objects and the full path was within the manipulator's workspace, after which the operator instructed the real manipulator to perform the same task by downloading the trajectory from IGRIP to the robot's controller. The physical robot then performed the same action.

3.6.3 Free-Form Object Recognition Method


3.6.3.1 Introduction

The QPSM method worked well on objects with fairly simple geometries (e.g., boxes, vessels, and pipes) which in itself is a useful result because, while seemingly complex, many DOE facilities are collections of many simply shaped objects, albeit in a complicated arrangement. However, the method also has difficulties with occlusions, i.e., objects that are partially hidden from the sensor by other objects. One consequence of this symptom is that QPSM-based Artisan interprets a single pipe passing behind an I-beam, for example, as two different pipe segments, one on either side of the I-beam. In addition, the segmentation step is computationally expensive and recognition results, though obtained with very little operator interaction, came slower than we felt reasonable.

To cope with those shortcomings we developed a new approach, known as the Free-form Object Method (FORM), in Phase III. The FORM's underlying technique, known as geometric indexing, uses a novel description of surface geometry to make comparisons between any two sets of 3-D data points. A unique property of this algorithm is that it requires no knowledge of the transformation between sensor views. The representation comprises descriptive spin-images associated with oriented points on the surface of an object. Constructed using single point bases, spin-images are data level shape descriptions that are used for efficient matching of oriented points. Spin-images are correlated to establish point correspondences between two views from which a rigid transformation that aligns the views is calculated. The transformations are then refined and verified using a modified iterative closest point algorithm. The geometric indexing algorithm is described in detail in Appendix A.

3.6.3.2 FORM Operating Scenario

As with the QPSM method, an operator commands Artisan to take an image of the scene with a Perceptron Model 5000 laser rangefinder. To eliminate some of the noise in the range data, multiple range images are taken and the median range value for each pixel is returned. Once the range image is displayed the operator draws a rectangular region of interest around the objects to be modeled. This reduces the data to be processed to that which is important to the task at hand. Figure 3.9 shows range and intensity images acquired with the Perceptron and five different regions of interest (the labels are presented here only for the purposes of illustration).



Raw image produced from laser scanner. The five different regions of interest that are selected are
A: Control Box, B: Elbow Joint, C: Fan, D: Small Valve, E: T-joint

Figure 3.9 Range data used in the FORM recognition example

Phase II work with amplitude-modulated scanning laser rangefinders (such as the Perceptrons) noted certain problems with their intrinsic characteristics particularly in regards to their behavior when measuring distances to certain surface materials and at high incidence angles. As part of the development of the Artisan system we undertook an experimental characterization of the 3-D laser rangefinder used for of our work and used those results to improve the data sets fed into the

object recognition algorithm. For example, as the incidence angle (measured between target surface normal and viewing direction) to a shiny, cylindrical objects increases beyond 20 degrees, range measurements become unreliable. Thus an additional step, after temporal and spatial filtering, in which vertices whose incidence angle exceeding a fixed threshold are eliminated from the raw range data prior to formation of the Cartesian mesh, was employed. This procedure also eliminates mixed pixels and mesh edges that cross range discontinuities. A practical implication of this result is that targets such as unpainted metal pipes frequently found in industrial facilities can be difficult to model correctly since the system has less good quality data to use.

A Cartesian surface mesh is then created from the region of interest by making each pixel a vertex and connecting vertices across rows and columns. The final step in scene data processing is the application of a low-pass filter to the mesh [21]. This filter smooths without shrinking, so it removes spurious noise from the mesh while still preserving the shape of objects. Figure 3.10 shows a shaded view of a scene surface mesh from the range image shown in Figure 3.9 before and after application of the filter.




Figure 3.10 Cartesian mesh before and after 3-D smoothing filter

Prior to object recognition, models of objects that are to be recognized must be created. Since we are interested in modeling complex scenes, a representation that accommodates arbitrary geometric shapes as chosen. A polygonal surface mesh is an established way to describe the shape of complex objects in computer graphics and it is amenable to our recognition algorithm which requires a surface represented as oriented points (3-D points with associated surface normals). The vertices of a surface mesh correspond to points on the surface of the object, and the normals at the vertices can be calculated by fitting a plane to the vertex and all of the neighboring vertices in the mesh. Using surface meshes places very few restrictions on the shape of objects that can be represented, making our recognition system extremely flexible.

An important requirement of any recognition system is the ability to generate the model representations used in recognition with relatively little effort. Given a CAD drawing of an object to be recognized, surface mesh generation is simple. The CAD drawing is imported into a CAD

package with finite element capabilities (e.g., Pro/ENGINEER). The finite element software is then used to automatically tessellate the surface of the object into triangular faces given some user defined constraints on minimum and maximum edge lengths. Figure 3.11 shows an example of a CAD model of a fan transformed into a surface mesh and Table 3.5 shows the surface mesh models created from CAD drawings used in Artisan's model library.









Figure 3.11 Steps to add a model to Artisan's library

Table 3.5 Objects in Artisan's library

Icon	ID #	Object	CAD (ProE)	Segmented	Tgrip	Spin Stack
	1001	10 cm pipe		X	X	X
	1002	20 cm pipe		X	X	X
	1003	30 cm pipe		X	X	X
	1004	7.5 cm pipe		X		X
	1005	15.0 cm pipe		X		X
	2001	7.5 cm x 10.0 cm T-joint		X	X	X
	2002	15 cm elbow joint 90 degree	X	X	X	X
	2003	7.5 cm x 7.5 cm 90 degree joint	X	X	X	X
	2004	7.5 cm x 7.5 cm 45 degree elbow joint	X	X	X	X
	2005	7.5 cm x 7.5 cm T-joint	X	X	X	X

Table 3.5 Objects in Artisan's library

Icon	ID #	Object	CAD (ProE)	Segmented	Tgrip	Spin Stack
	6001	graphite brick 30.5 cm x 10.2 cm x 10.2 cm				X
	7001	fan	X	X	X	X
	7002	large valve and T-joint	X	X	X	X
	7003	small valve	X	X	X	X
	7004	electrical control box	X	X	X	X

The representation we use for recognition assumes that the vertices in the surface mesh adequately describe the shape of the object. Implicit in this is the assumption that the vertices of the surface mesh are evenly distributed over the surface of the object. Occasionally, the model surface mesh generation process (FEM or Multi-view) will distribute vertices be too coarse or too fine for recognition. To solve this problem we use a mesh regularization algorithm that normalizes the lengths of edges in a mesh while preserving the shape of the object represented [13]. The end result is a surface mesh representation of appropriate resolution with an even distribution of points over the surface from which our recognition representation can be made.

Artisan recognizes objects by establishing correspondences between oriented points (surface mesh vertices with normals) on the surface of a model and oriented points in the scene. Every oriented point on an object has an associated spin-image that describes the shape of the object with respect to the oriented point. A spin image for an oriented point on a surface mesh is created as follows. For all other points on the surface mesh, the 2-D oriented point coordinates (α , the distance between the surface normal and the projection of the point into the tangent plan, and β , the distance between the point and tangent plane, as shown in the left side of Figure 3.12) are calculated. Next, the pixel in the spin image that is indexed by the oriented point coordinates is incremented by one. Some examples of spin images for a CAD surface mesh model are given in Figure 3.12.




Figure 3.12 Spin image basis geometry and example spin images

Matching points based on spin-images lies somewhere between geometric hashing [15] and structural indexing [19], two forms of model based recognition. Spin images of corresponding oriented points will be similar because the construction of spin images is based only on the shape of the object, and not the pose of the object. Therefore, a comparison of spin-images suffices to establish oriented point correspondence.


The models in the model library are processed for recognition, prior to task space modeling. For each model, the spin-images for each oriented point are created using the points on the model. These images are stored in a stack for later use in recognition. Processing the models off-line saves time during recognition.

After selection of the region of interest, the operator selects a model(s) from the model library to be recognized and localized in the scene. Spin-images from a random selection of oriented scene points are created and compared to all of the spin-images of the selected models. Point correspondences are established between model and scene oriented points whose spin images are similar. Clutter in the scene and symmetries in the models may cause incorrect correspondences, so correspondences are filtered by grouping them into sets that are geometrically consistent. From these sets of correspondences, plausible rigid transformations are calculated that align the model(s) with the scene data.


The plausible transformations are verified and refined with a modified iterative closest point algorithm [1]. The verification algorithm iteratively spreads point correspondences from those already established by spin-image comparisons over the surface of the model and the scene. At each iteration it recalculates the transformation from model to scene. If the model and scene are well aligned and have similar shape then the number of correspondences created will increase dramatically. Models that have a large fraction of their points in correspondence with scene data are reported as recognized in the scene.

Recognition results are presented to the operator in a 3-D viewer window that provides rotation, translation and zoom capabilities as well as several modes of rendering. Figure 3.13 shows wire frame representations of the recognized models for the five regions of interest selected in Figure 3.9 along with the original scene data shown as shaded surfaces. The user makes the


determination of whether the recognized objects are localized well enough to be inserted into the model of the robot's world. This confirmation step affords a final human check on the validity of the object recognition. Figure 3.14 shows the objects of Figure 3.13 as they appear in the scene relative to our Cincinnati Milacron T3 industrial manipulator.




Control Box



Elbow Joint



Fan



Small Valve



T-joint

Figure 3.13 Objects recognized using FORM overlaid on Cartesian meshes




Figure 3.14 TeleGRIP workcell showing objects recognized using FORM

4. Conclusions

We have gained an appreciation for the diversity of D&D tasks planned and proposed to transition many DOE facilities to a state in which they can be reused with unrestricted access. Outright demolition is an unacceptable solution to facility decommissioning due to risk of radioactive release, loss of recyclable material, and introduction of huge volumes of contaminated spoils into the environment. D&D of most facilities is expected to be a controlled de-construction that proceeds "inside out" in stages beginning with dismantling, decontamination of equipment and piping, sorting of wastes, size reduction and packaging, and ending with decontamination of the building structure itself.

Given the wide range of tasks that must be executed, the weak structure of the settings in which they are to occur, and the diverse environmental hazards within the DOE, it is unlikely that a single worksystem design can meet all the needs. However, despite a huge D&D agenda, the development of several worksystems, each one specialized to a particular task, cannot be justified at this time due to the lack of experience in D&D operations, the lack of specific regulatory requirements, and the lack of detailed plans for facility transition. The most reasonable way to proceed is, thus, the development of versatile worksystems that are able to handle a variety of tools and perform a wide range of tasks. These worksystems must combine brute force for heavy work with dexterity for fine manipulation; they must be reliable for extended use and adaptable to a range of work conditions and settings.

Equally important to system designs are the interfaces and telerobotic controllers that allow human operators to make effective use of remote worksystems to cope with unforeseen circumstances and react to changing situations. Since it is highly probable that D&D scenarios will always require a man-in-the-loop, the utility of the remote equipment will be largely dependent upon how well the man and machine are able to operate as combined system. This ability to work cooperatively is determined by the fidelity of interfaces between the worksystems and human operators and the extent to which the telerobotic controllers can assume responsibility for safe execution of tasks.

4.1 Mobile Worksystem Characteristics

4.1.1 Configuration

Mobile worksystems of Rosie and RWV scale have the maneuverability, high reach capability, and ability to handle, position and power a wide range of tools to accomplish D&D work in DOE defense nuclear facilities. Despite their relative compactness, they have significant reach and lifting capacity and are quite stable even with booms at full extension. Tethering is justified, since it allows unlimited delivery of primary electrical power and makes communications very reliable. Their modular design provides responsiveness to a work agenda that is broad and dynamic, though increased modularity can increase their relevance to a larger number of sites. For example, substituting onboard power generation or storage and wireless communications for the tether will extend the worksystem's operating range and make it possible to operate in settings where tethering is not viable.

With on-board tether management, tethering does not significantly hinder mobility. Electrohydraulic actuation has significant advantages over all-electric designs for this class of work machine. Primary motions (driving, steering, boom, manipulator, tether reel, etc.) need to be

forceful, but generally run at low speeds, so hydraulic actuation is appropriate. The availability of on-board hydraulic power also facilitates the use of many tools. One drawback of the RWV and Rosie are their large number of spaces, particularly between exposed hoses and cables, that could collect contaminants. Design for decontamination can always be improved.

4.1.2 Locomotion

Both the RWV and Rosie have mobility characteristics that are well-suited for work in flat floor environments. Two important capabilities, omni-directional locomotion and multi-mode steering, allow multiple modes of driving, of which curvilinear (also called "four-wheel steer" mode) and translational (also called "crab" mode), are used most often. Lateral driving, i.e., side to side, is particularly useful for tasks relative to walls and long runs of horizontal pipes. The third mode ("rotate about point") has utility, though fewer tasks require it. A simple and quick way to switch between modes, such as the push-button selector employed in both consoles, is absolutely essential.

Driving is a necessary aspect of every task that a mobile worksystem performs and can represent a significant fraction of the overall task execution time. Since remote driving is often more difficult than remote manipulation, better feedback to the operator can improve this aspect of mobile worksystem operations. A difficulty in remotely driving the RWV is judging its proximity to objects in its surroundings. All operators report that the main cameras provide the most useful perspectives for driving, but the pan and tilt positioning actuators move too slowly. To overcome this, operators often use the camera on the manipulator's wrist during driving because it can see the front and sides of the robot and can be rapidly positioned. Cameras dedicated to driving, particularly for assessing side clearances, are warranted.

It also appears advisable to incorporate navigation and automatic guidance capabilities. Position information could either be provided to the operator as a navigational indicator or be incorporated into automatic driving controls. Both would assist in maneuvers such as shuttling back and forth between two locations or moving along paths that had been driven previously. Some tasks, such as radiological mapping, will require knowledge of position in a global coordinate frame so that acquired measurements can be related to the locations from which they are collected. There are several mature technologies available for navigation of interior spaces, ranging from dead-reckoning (which works very well on flat floors) to more sophisticated guidance systems based on beacons, to correlation of observed features in the robot's environment with their known locations on a map.

4.1.3 Control

To effectively control the numerous worksystem motions and monitor its status, the mobile worksystem should have an on-board real time controller. Not counting light duty manipulators (the ones used in this program had their own control systems), both Rosie and RWV have over thirty controllable motions, many of which have to be coordinated. Controlling this many motions from an off-board location would have required a very large number of conductors between the robot and controller, so placing the controller onboard the robot was warranted. The division of control responsibilities between the onboard computers, which performed all time critical tasks such as closing servo loops and coordinating motions, and the console computer, which converted operator inputs into robot commands, also worked well.

Computer boards compatible with the software and operating system used to develop the RWV controller are presently available with radiation tolerances in excess of 10^6 rad. It is important that the controller measure and report the motion of all joints both for safety and for displaying the robot's configuration in the graphical simulation. Software should also be modularly designed and compatible with other robot controllers so that future upgrades can be made with minimum effort.

4.1.4 Manipulation

Manipulators that are roughly the size of, but substantially stronger than, the human arm are appropriate for remote D&D work. D&D tasks require manipulators that have at least six degrees of freedom; additional worksystem motions are needed to locate the base of the manipulator in the workspace.

Both manipulators used in this work (a Kraft GRIPS™ in Phase I and a borrowed Schilling Titan II in Phase III) have sufficient dexterity for the class of D&D tasks that were attempted. Operators quickly learn how to combine boom motions with manipulator motions to execute tasks, but generally favor the smaller manipulators, which are faster and more dexterous than the heavy manipulators. Remote manipulation tasks such as grasping and positioning tools require a great deal of dexterity, since they cannot be pre-planned in enough detail to pre-program a robot to execute them. Instead, operators use finesse to adjust their actions as a task proceeds. In many respects, the remote manipulator is a projection of the operator's arm in the remote working environment.

As shown in Table 4.1, mobile worksystems have more than the six degrees of freedom (DOFs) needed for general posing. When teleoperating the worksystem, these DOFs are used serially, i.e., the heavy manipulator is positioned with locomotion, then the tool deployment mechanism is positioned with heavy manipulator motions.

Table 4.1 Degrees of Freedom for several mobile worksystem configurations

base platform	tool deployment mechanism	Degrees of Freedom			
		locomotion	heavy manipulation	fine manipulation	total
RWV	Badger heavy tool positioner	3	3	2 ^a	8
RWV	manipulator	3	3	6	12
Rosie	single manipulator	3	4	6	13
Rosie	Dual Arm Work Module	3	4	9 ^b	16

a. one of the two heavy tools that the Badger can deploy

b. one of the two manipulators (6 DOFs) and the associated positioners on the DAWM base (3 DOFs)

Separating manipulator control from the rest of the robot (locomotion and boom) allows straightforward control of end effectors. Both of the manipulators used in this program have self-contained controllers that are independent of the mobile worksystem's controller. In most manipulation tasks, the two control systems were used independently, a scheme that operators are able to learn quickly and use effectively. The ability to lock out selected degrees of freedom is particularly useful.

If, however, the whole worksystem is to be operated automatically, the light duty manipulator must have a computer controller that allows it to be controlled like a robot. Further, complex

control schemes are required to resolve redundancies among the degrees of freedom. A reasonable compromise is to preserve the serial nature of motions. For example, control of locomotion might be kept separate and distinct from all manipulation. In general, however, coordination of all motions is required.

Position control alone is sufficient for some, but not all D&D tasks. All operators reported that tasks like dropping an object into a container, grasping an object that was free to move (e.g., a section of pipe laying on the floor) and grasping when there was little concern for damaging the object were easy to perform. In contrast, tasks like turning a valve or cutting through a pipe with a saw or grinder were rather difficult. The difference between these two classes of tasks is that the latter require control of both force and position while the former can be successfully accomplished with position control alone. Performance and efficiency can be improved through more advanced control schemes in which forces at the end effector are measured and fed back to the manipulator control system and/or the operator. Force information can be used in several ways ranging from simple limiting of end effector forces/torques in selected directions to elaborate compliant motion control schemes.

Better feedback to the operator can improve remote manipulation. Operators, particularly novices, often report difficulties in judging distances during manipulation tasks. In some manipulation tasks, operators reported difficulty achieving a desired end effector orientation while avoiding collisions between other parts of the robot and its surroundings. Difficult manipulation tasks can be facilitated by using task space scene analysis to model the manipulator's workspace, then employing automatic control to execute the task. Information about the manipulator's proximity to objects should be included a 3D display to help the operator relate to the scene.

4.1.5 Tooling

There is a wide variety of shelf-available tools that can be easily and inexpensively adapted for use with mobile worksystems; hydraulic tools are preferred for many D&D tasks. Advantages of hydraulic tools over electric counterparts include higher power densities, higher force/torque output, and more forgiveness in overload conditions. Simple modifications to shelf-available tools can simplify remote work. For example, the task of cutting pipes and placing the cut sections in a container could be performed more readily by adding a mechanism to hold the pipe as it is cut (so that the cut section does not fall to the floor) and by providing the means to carry the container around (so the robot does not have to drive back and forth to the container).

The hydraulic shears and grinder were both found to very effective cutting implements; the performance of electric saw, however, was rather disappointing. The shears were superior for cutting pipes, though grinder was somewhat more versatile (cutting pipe brackets, for example).

The ability to carry multiple tools each time the mobile worksystem is deployed is a key to remote D&D. Given anticipated operating scenarios in which mobile worksystems will perform surgical equipment dismantlement and highly selective decontamination, this approach appears to be preferred over carrying a single tool into the workspace. However, there may be situations in which the latter approach is justified by the volume of work that can be performed with a single tool, e.g., cutting a large number of pipes. Ultimately, the selection and number of tools will be driven by the amount and nature of the work agenda. Versatility is one of a mobile worksystem's strong suits, and every effort should be made to capitalize on this strength. Better schemes for tool

interchange are needed; use of a single standardized tool interface providing mechanical, electrical and hydraulic connections that can be made and broken remotely and automatically would facilitate this. Operation of many tools requires dexterity and precise manipulator control.

The Badger heavy tool positioning mechanism was a useful addition to the RWV since it allowed use of the tools too heavy for the Kraft manipulator and provides a mounting location for a tool holster. However, since the Badger lacks a roll motion, it is somewhat limited in the sense that orientation control lacks a degree of freedom. Likewise, Rosie would benefit from the addition of a yaw motion at the wrist of the heavy manipulator. One drawback of Rosie is the need to add a pressure intensifier and related components to run high pressure tools such as a hydraulic shears. A separate high pressure circuit should be incorporated in future designs.

4.1.6 Operator Interface

Because the human operator is essential for effective worksystem operation, the mobile worksystem's operator interface is a critical element of the overall man/machine system. Just as a well designed interface will facilitate remote D&D activities; a poorly designed interface can degrade the performance of an otherwise capable mobile worksystem. Attention to ergonomics and operator comfort is essential to achieving this goal. The operator interface design should emphasize the aspects of worksystem control that are used most often without comprising access to secondary functions. The approach used in the RWV and Rosie consoles – a small number of reconfigurable hard controls (joysticks and switches) for primary motions and a touchscreen with a menu of soft controls to operate less frequently used functions and to configure the hard controls – works well.

4.2 Task Space Scene Analysis Systems

Our TSSA current instantiation Artisan fulfills most of general requirements we had formulated and fills an identified technology gap. However, there are some deficiencies which should be addressed in future work.

Early in our development we experienced uncertainties in object model locations on the order of 10-15 cm. Through a sensitivity analysis, we determined that the range sensor was the predominant source of this error, which is not surprising given the range accuracy of the device we were using at the time. Substituting the ORNL scanner, which produces range errors on the order of 1-2 cm has improved model uncertainty commensurately. We are confident that the overall uncertainty can be substantially reduced by substitution of a more accurate sensor. Sub-centimeter range resolution is desirable but should be achievable.

Artisan-generated models are currently limited to objects visible in a single image with a 60° x 60° field of view which satisfies the requirements for modeled scene height and width. Model depth, however, is limited because task space scenes contain many objects stacked behind one another and only those closest to the sensor can be modeled accurately and completely. The solution to this problem is to move the sensor into multiple viewing positions and merge the data acquired from each perspective into a larger composite model. Though techniques to do that are available, they were developed for isolated objects and have not been applied to the types of scenes that D&D robots will encounter.

The geometric indexing approach to object recognition offers substantial improvements in performance over previous systems. First, it does not make any assumptions on sensor used for acquiring the range data. In fact, we have demonstrated the system both with an imaging laser rangefinder and with a light stripe rangefinder. By contrast, systems based on interactive stereo or line segmentation make strong assumptions on the geometry of the sensor and do not generalize nearly as well as the geometric indexing algorithm at the heart of Artisan.

Another drawback of competing techniques is that they tend to degrade the input range data by forcing it into one of the primitive classes. Our approach is general in that, by working directly with surface meshes, it does not make assumptions of the shape of the objects that can be recognized. Our earlier system restricted recognition to objects composed mostly of planar and quadric surfaces. This relegated Artisan to use in environments comprised only of objects with simple geometries, such as cylinders and boxes. Further, our earlier system's ability to recognize was much more dependent on the amount of range data available for a given object in the scene. This made it more susceptible to the occlusions that are prevalent in industrial process complexes; the earlier Artisan would not "understand" that two segments of a horizontal pipe laying behind a vertical pipe were actually the same object. In addition, the 3-D segmentation step is computationally expensive. The newer Artisan can do everything its predecessor did in about half the time.

Finally, building models for the recognition algorithms is straightforward. Any CAD system can be used as long as it can generate meshes, such as finite elements meshes. Although capability is not yet integrated in the Artisan system, it is also possible to generate models from range data directly. This would afford, for example, the capability to image a heretofore unknown object within a facility and recognize subsequent occurrences of it.

Artisan's ability to recognize objects that have complex geometry makes it an invaluable tool for modeling contaminated facilities that are served or acted upon by robots. Object recognition is preferred over space occupancy modeling or 3-D viewing since it provides semantic information in addition to geometry. This enables performance of robotic tasks on objects that are recognized. The geometric indexing algorithm Artisan employs has distinct advantages over alternative approaches since it can cope with partially occluded surfaces and changes in sensor viewpoint.

Future work will focus on expanding Artisan's model library, developing an automatic multi-view model generation system and enhancing Artisan's ability to recognize multiple objects simultaneously. We are working on a number of extensions to the object recognition algorithms described above. First, we will need to deal with large library of objects in real applications. This is not only because of the complexity of the environments, but also because allowing large sets of models will alleviate the burden on the operator to select models from the library. We are working on an approach in which all the spin images of the all the models are combined into a single stack. At recognition time, the spin image from the scene are compared to all the spin images in the combined stack. Because this approach is computationally expensive algorithm for reducing the number of spin images representing a model. This algorithm involves computing the principal directions of the space spanned by the spin images and to retain only the salient ones, thus reducing the number of spin images.

One weakness of the current recognition algorithm is that it takes more time to recognize a symmetrical object because of there are many different ways of pairing the spin images. The

technique outlined above for multiple objects will also work for identifying symmetries in objects and reducing the recognition time. Another limitation is that the object models are represented at a single scale. For example, we can recognize a pipe and its position in the image but we cannot compute its radius if it deviates too much from the model's radius. One approach is to parameterize the spin images as function of the scale of the objects. We are investigating this possibility.

Finally, a practical improvement to Artisan is the automatic selection of recognition parameters. Currently, recognition parameters such as mesh resolution or thresholds, are stored manually with each model. In fact, they can be computed directly from the models. We plan on incorporating this capability in Artisan.

5. Recommendations

5.1 Future Modifications to Rosie

(Part of this program included development of a modified version of Rosie (Rosie-C) for use in the CP-5 reactor dismantlement. The following items were discussed during the formal design review of Rosie-C.)

5.1.1 Decontamination

The hydraulic fluid in the system is cooled continuously by a forced convection fluid to air heat exchanger mounted on the robot. This heat exchanger exhausts air at a rate of about 2,500 SCFM straight out from the side of the robot, approximately 2 feet above the floor. This air flow could stir up contaminated particles on the floor of a work site. Exhausting the air in a different direction, such as straight up through the deck of the locomotor, out the back, or through a snorkel-like vent could alleviate the problem.

The hydraulic hoses used on the Rosie system have a braided outer covering. The same type of hose is used everywhere on the system. It was selected for abrasion resistance, small bending radius, and high pressure capabilities in a wide range of sizes. However, the braiding presents potential locations for contaminated particles to collect and become trapped on the hose surface. Porous electrical cables can similarly collect contamination. Alternative hose and cable types could rectify this problem.

The onboard electrical enclosure is a completely sealed box that uses an air-to-air heat exchanger to remove heat from it. It might be advisable to filter this air before it enters the heat exchanger to keep its interior cleaner for a long period of time.

To facilitate decontamination of any surfaces on the robot, all surfaces should be sealed and smooth. This can be accomplished by painting with a hard surface paint, such as epoxy, or by sealing with a surface treatment, such as anodizing, for aluminum.

5.1.2 Operations

There may be times when two-way voice communications between the robot and the console would be desirable, e.g., when tools are being installed or maintenance is being performed. Reallocation of tether conductors and adding an outlet on the robot for a headset could enable such communications.

Load capacity of the heavy manipulator varies according to its configuration; hence it is conceivable that an operator might unknowingly try to lift too much weight, this could result in damage to the heavy manipulator or loss of stability (tip over) of the locomotor. The former can be avoided with the simple addition of a load cell to the lifting hook and pressure relief valves on the shoulder actuators. Guarding against tip over is more complex, though addition of a tilt sensor to the locomotor can provide the operator with a simple warning at the console. Spare Parts Costs

Cameras near the boom tip can get dirty during tooling operations. A remote cleaning system, using either water or air, could circumvent this, but would require a local source of fluid on the robot. Simple "windshield wipers" is another alternative.

5.1.3 Alternative Configurations

Similar, scaled (both larger and smaller) mobile worksystems should be developed. No single machine can meet all of the DOE D&D needs. Through discussions with potential customers inside and outside the DOE community, we have already heard people say, "That robot is much too big for my application. I need something that will fit through a standard doorway." On the other hand we have also heard, "I need more reach and payload capability than the 26 feet and 1800 lb that Rosie offers." Therefore, we believe there is a need for scaled versions of Rosie to handle both bigger and smaller tasks than Rosie.

An untethered version of Rosie should be developed. Tethering is advantageous from the standpoints of power density, reliable communications, and elimination of exhaust from an internal combustion engine. However, it limits mobility in terms of working range and complexity of driving paths that can be achieved. There are situations where an untethered machine would be more appropriate, if not required. We have considered this during the design of the system. Rosie could be adapted to run on a conventional gas or diesel engine or even a propane gas engine. This would require some redesign since it would be preferable to run the hydraulic system directly from the engine, instead of with an electric motor, but this will be a required configuration in certain facilities.

An all-electric worksystem should be developed. In certain situations it is infeasible to operate a hydraulic system; where significant quantities of nuclear material are present in a facility, a hydraulic leak could mobilize the material and/or make for a criticality concern. Therefore, there is a need to provide a system without the hydraulic fluid required for Rosie. This would mean that all hydraulic actuators would need to be replaced with electro-mechanical devices. While it is doubtful that a fully electric machine could provide the same reach and payload characteristics as Rosie, this machine would still be of significant value in a number of applications.

5.2 Future Modifications to Artisan

5.2.1 Recognition of Individual Objects

Considerable improvement in the present Artisan system's ability to recognize individual objects can be made without fundamental changes to the algorithms. Most importantly, the precision of object pose and dimension estimates made by Artisan can be increased and the number of false matches reduced. These require enhancements to the recognition algorithms themselves. In addition, choice between either of the two recognition techniques, based on the task and the type of objects to be recognized, can be automated.

The internal representation used in the FORM (spin map stacks) can be statistically evaluated in order to compute the best parameters for object recognition. A histogram of correlation between the spin maps of a given model is currently used to determine the saliency of particular points of the model, and to help compute thresholds for recognition. A similar technique can be used to compute the saliency of the spin maps for various densities of surface points and to compute the optimal data resolution for a particular object. Such an algorithm would be used off-line during model generation. The expected results are improved accuracy and reduced processing time.

The quality of index-based registration is dependent on the number of points used on model and scene data sets, i.e., the resolution of the 3-D meshes used in the recognition, and the appropriate resolution varies depending upon the object. The resolution is currently set manually though it can

be computed automatically from the models. Furthermore, the current approach uses a reduced data set in the final registration step (the modified ICP algorithm). While increased data density improves registration accuracy, there is a trade-off between the improved accuracy and the increased execution time due to the larger number of points used in ICP. Strategies for choosing the appropriate data density based on a trade-off analysis need to be developed.

The QPSM allows for parameterization of object models, which is useful for constructing a task space scene model. For example, the scale factors along arbitrary axes can be recovered by the recognition algorithm and adjusted by the operator to stretch or shrink the object. However, the indexing approach does not allow such parameterization and needs to be extended to do so. This can be achieved by developing a means to measure the scale of an object, and scaling spin images appropriately. Alternatively, a method to compare spin images of different scale can be developed. Both approaches should be implemented and compared in terms of computational efficiency.

The two Artisan object recognition approaches perform differently on different classes of objects. For example, QPSM performs better on simple objects with full symmetries (e.g., pipes) and on objects that are almost-polyhedral, while FORM is better suited to recognize complex shapes like valves. Therefore, strategies for switching between recognition modes are needed for optimal use of Artisan. The main criteria for deciding among recognition strategies are symmetries, complexity of the object (as measured by the errors in the quadric and planar segmentations), and density of points required in FORM. Those criteria can be evaluated from the model mesh and combined into one or several metric used for deciding among recognition modes.

5.2.2 Recognition of Multiple Objects

An operational capability for simultaneous multi-object recognition could reduce the operator's effort required to construct the task space model and would make Artisan more applicable to tasks and scenarios where there are large numbers of objects of various types. The current Artisan requires the user to run the recognition algorithm for each model separately. An obvious improvement at the system level is to make it possible to recognize several objects in a single step. The user should be able to specify large sets of object models and to let the recognition algorithm determine which objects are present in the scene. Further, Artisan should be able to process several recognition operations concurrently in order to relieve the operator from searching through the database, selecting objects one by one

The geometric indexing approach was designed with simultaneous multi-object matching in mind. Multiple indexing tables (corresponding to multiple models being tested in a given run) can be grouped into a single table. This combined table can be used during the recognition step as-if it were a single model. After indexing, the best sets of point matches can be labeled as different objects to allow simultaneous identification of multiple objects in a scene. Encouraging preliminary results on multi-object matching have been obtained off-line using complex objects from the medical domain.

Those techniques can be expanded upon to realize simultaneous multi-object recognition. Fundamental issues, such as the structure of the combined table, the control algorithm for dividing sets of point matches into subsets corresponding to different objects, need to be addressed, as do operational issues such as controlling the size of the combined table and managing multiple recognition thresholds.

5.2.3 Use of Range Data

Range data can be used more effectively in scene analysis. Artisan should be able to self-determine if enough range data has been acquired and is of high enough resolution to create the required task scene model. Work scope in this area includes developing and testing techniques to merge data from multiple sensor views, to select the appropriate sensor viewpoints, and to optimize data density. Artisan presently uses only a single view from the range sensor in order to identify objects in a scene. In practice, the degree of occlusion of typical objects in a single view may render recognition unreliable (or in some cases, impossible). The alternative is to view the scene from different angles and to perform recognition using the combined data set. However, compliance and other inaccuracies in the sensor positioning mechanism often make estimation of sensor pose based on motion records unreliable. Therefore it is important to develop techniques that can compute the transformation from the data sets directly, without relying on exact knowledge of the sensor pose. Preliminary results have been obtained using Artisan's geometric indexing algorithm applied to laser and light striper range data. However, these were obtained in off-line processing and are not currently available to the Artisan user at run time.

In general, it may be necessary to combine more than two views. The simplest approach would be to register pairs of views sequentially and to combine the resulting transformations. It is well known that such an approach leads to error accumulation as the number of views increases, and that the resulting transformation depends on the order in which the views are matched. An alternative, for which a coded algorithm exists, is to compute an optimal estimate of the transformation using all the views at once. This technique requires initial matches between views, which can still be obtained using the pairwise registration of views. Once the transformations between data sets have been computed, all the data points can be transformed into a single reference frame, structured into a mesh, and processed with either Artisan object recognition algorithm.

The use of multiple views can be evaluated using the existing object recognition algorithms. In particular, the trade-off between the additional time spent in acquiring and registering images and the time spent in correcting poor recognition results due to occlusions will be analyzed. Metrics to be used in this evaluation will include: accuracy of the registered model, combined computation and image acquisition time, and percentage of recognition failures.

5.2.4 User Interface

Some immediate opportunities to enhance the usability of Artisan are evident. The development program that created the existing Artisan system emphasized fundamental research and algorithm development in lieu of rigorous human factors engineering of the user interface. Though Artisan is a fully functional TSSA capability, a significant level of knowledge about its underlying algorithms, the specifics of the software implementation, interpretation of range data, and general computer vision is required to use the system properly.

There are several opportunities to reduce the operator expertise required to run Artisan. First, since the average user is more familiar with remote viewing systems (essentially closed circuit television) than with range data, it is desirable to substitute the range data display with video imagery. This requires collocating a video camera with the range sensor and calibrating their fields of view such that the range images and video images are spatially registered to each other. Second, there are several parameters associated with each step in the processing that have to be

properly "tuned" to achieve good results. Experimentation during the course of Artisan development indicate relationships between these algorithm parameters and parameters that describe the task space (such as the type of objects that are to be recognized, the size of the objects, and nominal distance between the sensor and the targets). These relationships can be exploited to automate setting of parameter values, thus alleviating the operator of this duty. Finally, there is substantial room for improvement in the basic human factors aspects of the interface.

Other changes to the Artisan interface are required to fully implement the enhancements described above. For example, while the current Artisan interface requires the user to select each object independently by clicking on an icon, this becomes unacceptable if large groups of objects can simultaneously processed by the recognition algorithm. Ways to group models through pre-processing and ways to help the user to select "natural" groups of models need to be addressed. Brand new user interface elements are also needed to support the use of multiple views. In particular, while only a single view can be displayed at any one time in the current Artisan system, additional functionalities need to be developed to display multiple images, select regions of interest in multiple images, display combined meshes, and select sensor vantage points.

5.3 Integrated Systems & Demonstrations

More performance data has to be generated to prove the viability of remote worksystems for D&D. To prove Rosie's utility, we must demonstrate that D&D operations can be effectively and efficiently performed with remote equipment. Experiments need to be conducted under controlled conditions (i.e., in mock-ups) to acquire data about how long it takes to execute isolated D&D tasks remotely. Full scale scenarios must also be executed to understand inter-relationships of work tasks and issues related to deployment, recovery and maintenance. These tests will provide valuable information and begin to prove that robots can achieve reasonable throughput. Such proof is necessary in order for stakeholders to consider robots a valid D&D technology. In addition, more tooling needs to be adapted and demonstrated to show the breadth of potential for remote work.

Only by putting the worksystem into hot service will its true viability be shown. Our project, indeed the entire RTDP D&D robotics program, is aimed at doing real D&D work, so we must continue to identify real applications where Rosie can be applied. Then we must work to educate and inform operations people to convince them that Rosie is the right solution for their problems. Only through real, successful operations will we prove that Rosie is a robust technology. A great deal will be learned from the CP-5 experience. It is of utmost importance that this information be captured and disseminated.

The benefits of semi-automatic TSSA have to be measured rigorously. Artisan's real utility depends on the cost (the one-time capital investment in sensors and computing plus the operational cost of time spent recognizing objects before any work gets done) of creating 3-D geometric information and how it is subsequently used. An immediate opportunity is to create graphical displays that faithfully represent the robot and its workspace, thus providing the robot operator with a sophisticated aid for teleoperation. Our experience, like that of others in the DOE robotics community, indicates that such a capability is in itself quite useful as a means to increase operator awareness of the remote situation. It also allows the operator to practice difficult maneuvers in the forgiving environment of a robot simulation before attempting the real thing.

Much greater benefits are possible, however, when the 3-D information is accessed by the robot control system, admitting the use of automatic planners. The simple demonstrations we conducted is an early data point that verifies this hypothesis, but testing in a wider range of tasks and settings is needed to provide definitive proof.

6. References

- [1] P. Besl and N. McKay, "A Method for Registration of 3-D Shapes," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 14, no. 2, pp. 239-256, 1992.
- [2] C. Chua and R. Jarvis, "3-D Free-form Surface Registration and Object Recognition," *Int'l J. Computer Vision*, vol. 17, pp77-99, 1996.
- [3] J. Devore, *Probability and Statistics for Engineering and Sciences*, Brooks/Cole, Belmont, CA, 1987.
- [4] R. Duda and P. Hart, *Pattern Classification and Scene Analysis*, Wiley-Interscience, New York, 1973.
- [5] O. Faugeras and M. Hebert, "The Representation, Recognition and Locating of 3-D Objects," *Int'l J. Robotics Research*, vol. 5, no. 3, pp. 27-52, 1986.
- [6] O. Faugeras, *Three-Dimensional Computer Vision: A Geometric Viewpoint*, MIT Press, Cambridge, MA, 1993.
- [7] W.E.L. Grimson, *Object Recognition by Computer: The Role of Geometric Constraints*, MIT Press, Cambridge, MA, 1990.
- [8] A. Guéziec and N. Ayache, "Smoothing and Matching of 3-D Space Curves," *Int'l J. Computer Vision*, vol. 12, no. 1, pp. 79-104, 1994.
- [9] M. Hebert, J. Ponce, T. Boult and A. Gross, Eds., *Object Representation in Computer Vision*, Springer-Verlag, Berlin, 1995.
- [10] Y. Hecker and R. Bolle, "On Geometric Hashing and the Generalized Hough Transform," *IEEE Trans. Systems, Man and Cybernetics*, vol. 24, no. 9, pp. 1328-1338, 1994.
- [11] B.K.P. Horn, "Closed-form Solution of Absolute Orientation using Unit Quaternions," *J. Optical Soc. Amer.*, vol. 4, no. 4, pp. 629-642, 1987.
- [12] A. Johnson and M. Hebert, "Recognizing Objects by Matching Oriented Points," Carnegie Mellon Robotics Institute Tech. Report CMU-RI-TR-96-04, June 1996.
- [13] A. Johnson and M. Hebert, "Control of Mesh Resolution for 3-D Object Recognition," Carnegie Mellon Robotics Institute Tech. Report CMU-RI-TR-96-20, October 1996.
- [14] A. Johnson, P. Leger, R. Hoffman, M. Hebert, and J. Osborn, "3-D Object Modeling and Recognition for Telerobotic Manipulation," *Proc. Intelligent Robots and Systems 1995 (IROS '95)*, pp. 103-110.
- [15] Y. Lamdan and H. Wolfson, "Geometric Hashing: A General and Efficient Model-based Recognition Scheme," *Proc. Second Int'l Conf. Computer Vision (ICCV '88)*, pp. 238-249, 1988.
- [16] Y. Lamdan and H. Wolfson, "On the Error Analysis of 'Geometric Hashing'," *Proc. Computer Vision and Pattern Recognition 1991 (CVPR '91)*, pp. 22-27, 1991.
- [17] I. Rigoutsos and R. Hummel, "A Bayesian Approach to Model Matching with Geometric Hashing," *Computer Vision and Image Understanding*, vol 62, no. 1, pp 11-26, July 1995.

- [18] D. Simon, M. Hebert and T. Kanade, "Real-time 3-D Pose Estimation using a High-speed Range Sensor," *Proc. Int'l Conf. Robotics and Automation (R&A '94)*, pp. 2235-2241, 1994.
- [19] F. Stein and G. Medioni, "Structural Indexing: Efficient 3-D Object Recognition," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 14 no. 2, pp. 125-145, 1992.
- [20] R. Szeliski, D. Tonnensen and Dimitri Terzopoulos, "Modeling Surfaces of Arbitrary Topology with Dynamic Particles," *Proc. Computer Vision and Pattern Recognition 1993 (CVPR '93)*, pp. 82-87.
- [21] G. Taubin, "A Signal Processing Approach to Fair Surface Design," *Proc. Computer Graphics 1995 (SIGGRAPH '95)*, pp. 351-358, 1995.
- [22] Z. Zhang, "Iterative Point Matching for Registration of Free-form Curves and Surfaces," *Int'l J. Computer Vision*, vol. 13, no. 2, pp. 119-152, 1994.