Research Institute for Advanced Computer Science

ANNUAL REPORT

October 1998 through September 1999

Cooperative Agreement: NCC 2-1006

Submitted to:

Contracting Office Technical Monitor:
Anthony R. Gross
Associate Director for Programs
Information Sciences & Technology Directorate
NASA Ames Research Center
Mail Stop 200-6

Submitted by:
Research Institute for Advanced Computer Science (RIACS)

**An Institute of the
Universities Space Research Association (USRA)**

RIACS Principal Investigator
Dr. Barry M. Leiner
Director

_____
Dr. Barry M. Leiner

TABLE OF CONTENTS

I   Introduction and Overview

The Research Institute for Advanced Computer Science (RIACS) carries out basic research and technology development in computer science, in support of the National Aeronautics and Space Administration's missions. RIACS is located at the NASA Ames Research Center (ARC). It currently operates under a multiple year grant/cooperative agreement that began on October 1, 1997 and is up for renewal in the year 2002.

ARC has been designated NASA's Center of Excellence in Information Technology. In this capacity, ARC is charged with the responsibility to build an Information Technology Research Program that is preeminent within NASA. RIACS serves as a bridge between NASA ARC and the academic community, and RIACS scientists and visitors work in close collaboration with NASA scientists. RIACS has the additional goal of broadening the base of researchers in these areas of importance to the nation's space and aeronautics enterprises.

RIACS research focuses on the three cornerstones of information technology research necessary to meet the future challenges of NASA missions:

- Automated Reasoning for Autonomous Systems
  Techniques are being developed enabling spacecraft that will be self-guiding and self-correcting to the extent that they will require little or no human intervention. Such craft will be equipped to independently solve problems as they arise, and fulfill their missions with minimum direction from Earth.

- Human-Centered Computing
  Many NASA missions require synergy between humans and computers, with sophisticated computational aids amplifying human cognitive and perceptual abilities.

- High Performance Computing and Networking
  Advances in the performance of computing and networking continue to have major impact on a variety of NASA endeavors, ranging from modeling and simulation to data analysis of large datasets to collaborative engineering, planning and execution.

In addition, RIACS collaborates with NASA scientists to apply information technology research to a variety of NASA application domains. RIACS also engages in other activities, such as workshops, seminars, and visiting scientist programs, designed to encourage and facilitate collaboration between the university and NASA information technology research communities.

I.A   Summary of FY99 Activity

During the year October 1, 1998 through September 30, 1999, RIACS engaged in a number of research projects collaboratively with NASA scientists. Well over 100 publications and presentations resulted from this work. Section II describes each of those projects along with a list

of publications. In addition, RIACS ran a bi-weekly seminar series. The seminars are shown in Section III.

Staffing at RIACS grew considerably during the year, from 13 scientists at the beginning of the year to 20 scientists at the end. This continued the growth in the prior year (from 3 scientists to 13). In addition, 15 visiting scientists and 6 visiting students spent time at RIACS during the year. There were two changes in the Director during the year. Dr. Robert Moore left RIACS at the beginning of the year. Dr. Michael Raugh served as Interim Director while a search for a new Director was conducted. In June 1999, Dr. Barry Leiner was appointed Director. Section IV provides more detail on the staffing.

II   RIACS Projects During FY99

II.A   Automated Reasoning for Autonomous Systems

II.A.i   Remote Agent

II.A.i.a   Task Summary:

The main goal of this task is to support the continued development, implementation and demonstration of the Remote Agent. The Remote Agent is an autonomous software agent for controlling spacecraft and other complex systems. It provides three key capabilities; generating detailed plans from high-level goals, executing the resulting plans safely, and monitoring of the execution to assess and determine the actual state of the spacecraft. Additionally, this task aims at the further development of the technologies of the Remote Agent, including model-based diagnosis, and autonomous planning and scheduling.

II.A.i.b   RIACS Staff:

Ari K. Jonsson
P. Pandurang Nayak
Barney Pell

II.A.i.c   Overview:

Current spacecraft control technology relies heavily on a relatively large team of highly skilled mission operators to generate detailed time-ordered sequences of commands to step the spacecraft through each desired activity. Each such sequence is carefully constructed to ensure that all known operational constraints are satisfied. The autonomy of the spacecraft is limited.

The Remote Agent (RA) approach to spacecraft commanding and control is different from traditional operations. In the Remote Agent approach, more 'smarts' are put on board the spacecraft, as the operational rules and constraints are encoded in the flight software. The remote agent can autonomously achieve goals specified by missions operators, and can thus be thought of as an autonomous 'remote agent' of the operators. The operators do not know the exact conditions on the spacecraft, so they do not tell the agent exactly what to do at each instant of time. They do, however, tell the agent exactly which goals to achieve in a specified period of time.

The Remote Agent Experiment is part of the New Millennium Program's Deep Space One (DS1) mission. The New Millennium Program is designed to validate high-payoff, cutting-edge technologies to enable those technologies to become more broadly available for use on other NASA programs. The Remote Agent Experiment (RAX) has multiple objectives. A primary objective is to provide an on-board demonstration of spacecraft autonomy. This includes both nominal operations with goal-oriented commanding and closed-loop plan execution, and fault protection capabilities with failure diagnosis and recovery, on-board re-planning following unrecoverable failures, and system-level fault protection.

Other equally important, and complementary goals of the experiment are to decrease the risk (both real and perceived) in deploying RAs on future missions and to familiarize the spacecraft engineering community with the RA approach to spacecraft command and control. There are three parts to achieving this goal. First, a successful on-board demonstration requires the integration of the RA with the spacecraft flight software. This integration provides valuable information on required interfaces and performance characteristics, and alleviates the risk of carrying out such integration on future missions. It also serves to familiarize systems engineers and flight software engineers with the integration of RAs with traditional flight software. Second, a perceived risk of deploying RAs is related to its ability to synthesize new untested sequences in response to unexpected situations. This risk is addressed by demonstrating a layered testing methodology that serves to build confidence in the sequences synthesized by the RA, in a variety of nominal and off-nominal situations. Third, operating the experiment with close cooperation between RA team members and DS1 ground operators. This serves to familiarize the ground operations community with benefits and costs of operating a spacecraft equipped with an RA.

Three separate Artificial Intelligence technologies are integrated to form the Remote Agent: an on-board planner-scheduler, a robust multi-threaded executive, and Livingstone, a model-based fault diagnosis and recovery system. The planner-scheduler generates detailed action plans from the specified high-level goals, while taking into account safety and operations rules. The executive then executes the concurrent plans according to the capabilities of the spacecraft systems. Finally, the diagnosis and recovery system monitors the execution and determines the state of the spacecraft by comparing observed results to expected results.

II.A.i.d   Project Description:

The main component of this project is the support of the Remote Agent Experiment. To prepare for the experiment, a number of goals must be achieved. Among those are the delivery, testing and validation of the Remote Agent system. In parallel with these activities, the system must be moved from the development platform to the computation platform that is on board the spacecraft. Due to the limited availability of high-fidelity spacecraft testbeds, this process uses a sequence of testbeds, each more faithful to the spacecraft than the previous one. The experiment itself starts with the uploading of the completed Remote Agent to the Deep Space One spacecraft. The Remote Agent will then be activated and a number of scenarios run, in order to achieve the validation objectives. These include on-board planning, failure identification, and re-planning after a planned execution failure.

A component of this project, related to the Remote Agent experiment, is the development of ground tools to watch and analyze the telemetry data that is received from the spacecraft during the experiment. Since the Remote Agent is an autonomous system, the needed tools are much different from traditional telemetry analysis tools. This task supports the development of two different ground tools. The first, called Stanley, is a ground operations system for Livingstone. Based on the telemetry received from the spacecraft, the system determines the state and the actions of the on-board diagnosis system, and presents it in a graphical fashion. The other ground tool supports the on-board planner. Based on information received from the spacecraft telemetry, the tool emulates the planning process taking place on board the spacecraft and displays the results graphically.

Finally, a third component of this project is the continued development of technology from the Remote Agent. This includes further work on model-based mode identification, and the development of the next generation Remote Agent planner.

II.A.i.e   Accomplishments during FY99:

The primary focus of this task was on the DS1 Remote Agent Experiment (RAX), for which P. Pandurang Nayak was the deputy project lead. The culmination of that work was the successful completion of the experiment, on board the Deep Space One spacecraft, in May of 1999. That event has been hailed as one of the landmark accomplishments in artificial intelligence. In addition, the Remote Agent has been selected as one of two systems to receive NASA's Software of the Year award for this year.

The two key accomplishments for the Remote Agent were the completion of the testing/transition phase that took it from development to flight software, and, of course, the experiment itself.

The preparation work was primarily focused on testing the Remote Agent and moving it from the flight software testbed to the actual project testbeds. These project testbeds duplicate the actual spacecraft hardware to a much greater degree than the flight software testbeds. This process involved the transition of the Remote Agent to a variety of platforms, leading up to the main spacecraft testbed. The systems involved included the Sun/Solaris system used for development, the PowerPC/VxWorks system called "Babybed", the flight software testbed called "Papabed", the project testbed called "HotBench" and finally the high-fidelity project spacecraft testbed. As part of the preparation work, working closely with the DS1 team, a number of Operations Readiness Tests were designed and performed, to ensure that the experiment itself would go smoothly.

- The Remote Agent Experiment itself ran on the DS1 spacecraft in the week of May 17, 1999. It successfully demonstrated all validation criteria during this experiment. There were some unexpected glitches, but these turned out to allow further demonstration of the Remote Agent's ability to generate new mission plans at short notice.

- In support of the Remote Agent Experiment, two ground tools were developed, in collaboration with other Remote Agent team members, to analyze the progress of the experiment. One was Stanley, which determined and displayed the state of the spacecraft, based on information received from the on board mode identification and recovery system, called Livingstone. The other used telemetry information to emulate the expected progress of the on-board planner and display it graphically. The latter tool was the result of a collaborative project between our group and members of the TMOD group at JPL, aimed at connecting an existing missions scheduling tool (Apgen) with the Remote Agent Planner. Continuing the research thrusts that made the Remote Agent possible, further progress has been made on two of the key technology areas, model-based mode identification and autonomous planning.

- In the area of model-based mode identification, progress has been made on the development of approximation schemes for tracking the state of a spacecraft or another complex system.

The problem at hand is that the diagnosis system needs to be able to track multiple system trajectories. Unfortunately, in the worst case, there can be an exponential number of trajectories that need to be tracked. This is clearly infeasible. To address this issue, progress has been made on the development of various approaches to approximate the representation of these multiple states, and the analysis of their properties.

- Significant progress has also been made in research following up on the success of the Remote Agent planner. The result is the design and implementation of a next generation Remote Agent planning system. This system provides the same theoretical capabilities as the original planner, while providing important theoretical and practical advantages. Among those are the applicability of search techniques that are not based on chronological backtracking, and the use of an effective, general constraint reasoning mechanism.

II.A.i.f   Publications and Presentations during FY99:

1. Douglas Bernard, Gregory Dorais, Edward Gamble, Bob Kanefsky, James Kurien, Guy K. Man, William Millar, Nicola Muscettola, Pandurang Nayak, Kanna Rajan, Nicolas Rouquette, Benjamin Smith, William Taylor, and Yu-Wen Tung, "Spacecraft Autonomy Flight Experience: The DS1 Remote Agent Experiment", *Proceedings of the AIAA*, Albuquerque, New Mexico, 1999.

2. Ari Jonsson, Paul Morris, Nicola Muscettola and Kanna Rajan, "Next Generation Remote Agent Planner", in the *Proceedings of the 5th International Symposium on Artificial Intelligence, Robotics and Automation in Space (iSAIRAS-99)*.

3. P. Pandurang Nayak, Douglas E. Bernard, Gregory Dorais, Edward B. Gamble Jr., Bob Kanefsky, James Kurien, William Millar, Nicola Muscettola, Kanna Rajan, Nicolas Rouquette, Benjamin D. Smith, William Taylor and Yu-Wen Tung, "Validating the DS1 Remote Agent Experiment", in the *Proceedings of the 5th International Symposium on Artificial Intelligence, Robotics and Automation in Space (iSAIRAS-99)*.

4. Barney Pell, Gregory A. Dorais, Christian Plaunt and Richard Washington, "The Remote Agent Executive: Capabilities to Support Integrated Robotic Agents", submitted to *Autonomous Robotics Journal* for publication in 1999.

II.A.i.g   Future Plans:

The task has been completed with the successful experiment on board the Deep Space One spacecraft. However, the development of advanced planning and scheduling capabilities, building on the techniques developed for the Remote Agent, continues. More information on that task can be found in the "Advanced Planning and Scheduling" task (II.A.iii).

II.A.i.h   Problems Encountered:

None.

II.A.i.i   Problem Resolution:

None needed.

II.A.ii   <u>Automated Software Engineering</u>

II.A.ii.a   Task Summary

Research within the ASE task is primarily concerned with the development of advanced tools for automated software synthesis and automated software verification & validation (V&V).

II.A.ii.b   RIACS Staff:

Bernd Fischer
SeungJoon Park
Charles Pecheur
Willem Visser

II.A.ii.c   Overview:

Software is a major enabling technology at NASA. The primary goals of this task are to reduce the cost of developing this software and improve the reliability of the software produced.

Researchers and Visiting Scientists for the Automated Software Engineering task primarily develop advanced tools for automated software synthesis and automated software V&V. This entails a spectrum of activities from fundamental research, case studies with NASA software systems in conjunction with NASA software development organizations, development of new algorithms, and tool implementation.

II.A.ii.d   Project Description:

- Perform research to develop, design and implement techniques that will put verification tools into the hands of system designers.
- Investigate the application of model checking to find errors in software:
    - Verification of Autonomous systems
    - Verification of Avionics software
    - Program Verification
    - Abstraction techniques to reduce software complexity for model checking.
- Advanced research and development on software synthesis tools using deductive synthesis and other techniques from formal methods.
    - Synthesis of Data analysis algorithms
- Develop techniques to allow machine assistance for software reuse.

II.A.ii.e   Accomplishments during FY99:

- Explored the application of model checking to verify models for the HSTS planner from the Automated Systems group (ASG, code IC).
- Verification of the Livingstone System (ASG).

- A translation program from Livingstone models into SMV's input language.
- Support of Livingstone users at Kennedy Space Center (KSC) in getting acquainted with the principles of formal V&V and using the translator in KSC applications.
- Developed a formal definition of Livingstone models, to be used as a reference for discussing and asserting the correctness of the verification software.
- Verification of the SAFE protocol, developed at NASA Goddard for automatic data transfer between space station and ground stations via an unreliable intermittent connection, using the Java PathFinder tool developed by the ASE group.
- Developed a new abstraction algorithm, which automatically converts a concrete program to a corresponding abstract program with respect to user provided abstraction criteria, thus enabling larger programs to be analyzed by model checking.
- Found a time-partitioning error in the DEOS real-time operating system used by Honeywell to schedule applications in certain business aircraft. This error was not uncovered during testing and we developed a systematic approach to allow model checking to find the error.
- Extended the SPIN model checker to allow active objects to be used in modeling a system. This allowed more efficient translation from C++ and Java programs to SPIN for model checking.
- Developed a prototype Java model checker that is, to the best of our knowledge, the first model checker to handle the floating-point numbers and garbage collection.
- Built a prototype system, which applies deductive technology to synthesize efficient data analysis programs (e.g., classification, time series and image analysis) from high-level specifications in an extended Bayesian Network notation. Applied the prototype successfully to a series of standard problems; current work is on shadowing the data analysis tasks for the Vulcan-project (finding extra-solar planets using a photometric approach).

II.A.ii.f   Publications and Presentations during FY99:

1. J. Penix, B. Fischer, J. Whittle, G. Pour, J. van Baalen: "Automating Component Integration for Web-based Data Analysis Applications", *IEEE Aerospace Conference,* Big Sky, MT, 2000.

2. B. Fischer, J. Whittle: "An Integration of Deductive Retrieval into Deductive Synthesis"*, IEEE Automated Software Engineering Conference (ASE-99)*, Cocoa Beach, FL, October 1999.

3. B. Fischer: "Specification-Based Browsing of Software Component Libraries", *Journal of Automated Software Engineering*, To Appear, 1999 (Also available as RIACS technical report).

4. T. Baar, B. Fischer, D. Fuchs*:* "Integrating Deduction Techniques in a Software Reuse Application"*, Journal of Universal Computer Science*, Volume 5, No. 3, 1999, pp. 52-72.

5. W. Buntine, B. Fischer, T. Pressburger: "Towards Automated Synthesis of Data Mining *Proc. ACM International Conference on Knowledge Discovery & Data Mining*, San Diego, CA, August 1999.

6. B. Fischer, M. Lowry, J. Penix: "Intelligent Component Retrieval via Automated Reasoning", *Proc. AAAI-99-Workshop on Intelligent Software Engineering*, Orlando, July 1999.

7. W. Buntine, B. Fischer, K. Havelund, M. Lowry, T. Pressburger, S. Roach, P. Robinson, J. van Baalen: "Transformation Systems at NASA Ames", *Proc. ICSE-21-Workshop on Software Transformation Systems (STS'99)*, Los Angeles, May 1999, pp. 8-13.

8. B. Fischer, M. Lowry, J. Penix: "Tradeoffs in Behavioral Descriptions: Automated Reasoning for Dynamic Assembly of Components", Position paper for the *DARPA-Workshop on Software Behavior Description*, St. Thomas, U.S. Virgin Islands, December 1998.

9. B. Fischer: "Specification-Based Browsing of Software Component Libraries", *Proc. IEEE Automated Software Engineering Conference (ASE-98)*, Honolulu, HI, October 1998, pp. 246-254. Received Best Paper Award.

10. W. Visser and H. Barringer: "CTL* Model Checking for SPIN", *Proceedings of the 4th SPIN Workshop*, Paris, France. November 1998. This paper was selected for publication in the journal Software Tools for Technology Transfer (to appear).

11. W. Visser, Klaus Havelund and John Penix. *"Adding Active Objects to SPIN", Proceedings of the 5th SPIN Workshop*, Trento, Italy, July 1999.

12. J Penix, W. Visser, E. Engstrom, A. Larson and N. Weininger: "Verification of Time Partitioning in the DEOS real-time Scheduling Kernel", Submitted for publication.

13. J. Penix, C. Pecheur and K. Havelund: "Using Model Checking to Validate AI Planner *Proceedings of the 23rd Annual Software Engineering Workshop*, NASA Goddard, December 1998.

14. R. Gross, K. R. Sridhar, W. E. Larson, D. J. Clancy, C. Pecheur, and G. A. Briggs: "Information Technology and Control Needs for In-Situ Resource Utilization", In *Proceedings of the 50th IAF Congress*, Amsterdam, 4-8 October 1999.

15. C. Pecheur: "Advanced Modelling and Verification Techniques Applied to a Cluster File , *IEEE Automated Software Engineering Conference (ASE99)*, Cocoa Beach, October 1999.

16. Clancy, W. Larson, C. Pecheur, P. Engrand and C. Goodrich: "Autonomous Control of an in-Situ Propellant Production Plant", *Technology 2009 National Conference*, Miami Beach, FL, 1-3 November 1999.

17. S. Park and D. Dill: "An Executable Specification and Verifier for Relaxed Memory Order", *IEEE Transactions on Computers*, Vol-48 (2), pages 227-235, February 1999.

18. S. Das, D. Dill and S. Park: "Experience with Predicate Abstraction", *Proceedings of 11th International Conference on Computer Aided Verification*, Trento, Italy, July 1999.

Presentations:

1. Charles Pecheur, "Towards Formal Verification of Autonomy Software", *2nd Assurance Technology Conference*, NASA Glenn Research Center, 25-26 May, 1999.

2. Charles Pecheur, Peter Engrand, "V&V of MBAS", Presentation, tutorial and demonstration at Kennedy Space Center, 18 August 1999.

3. Charles Pecheur, "Formal Verification of Autonomous Systems", Florida Institute of Technology, 15 October 1999.

4. Bernd Fischer, "An Integration of Deductive Retrieval into Deductive Synthesis", *Presentation at ASE-99*, 13 October 1999.

5. Bernd Fischer, "Specification-Based Browsing of Software Component Libraries", Invited talk at Rice University, Houston, TX, 11 January 1999.

6. Bernd Fischer. "Specification-Based Browsing of Software Component Libraries", *Presentation at ASE-98*, 14 October 1998.

7. Bernd Fischer, "Evaluation of Automated Software Engineering Systems", Invited Panelist, *ASE99 Conference*, 15 October 1999.

8. Willem Visser, "CTL* Model Checking for SPIN". *Presentation at SPIN98*, Paris, 2 November 1998.

9. Willem Visser, "Software Model Checking at NASA Ames", Invited talk at Kansas State University, 5 April 1999.

10. Willem Visser, John Penix, "DEOS Verification", Invited talk at Honeywell HTC, Minneapolis, MN, 7 April 1999.

11. Willem Visser, "Adding Active Objects to SPIN", *Presentation at SPIN99*, Trento, Italy, 12 July 1999.

12. Willem Visser, "The Future of SPIN", Invited Panelist at *SPIN99 Workshop*, Trento, Italy, 12 July 1999.

13. Willem Visser, "Will Model Checking Scale to Software?", Invited Panelist, *ASE99 Conference*, 14 October 1999.

II.A.ii.g   Future Plans:

- Develop the general framework for the Generic Verification Environment (GVE) that will be used to find errors in software descriptions given in a wide variety of software notations, including design and implementation formalisms.
- Develop the Java and Livingstone Virtual Machines to be used within the GVE in order to do V&V of Java programs and the Livingstone fault recovery system, respectively.
- Continue research on the V&V of autonomous systems, specifically the work on the V&V of Livingstone and ISPP domain models.
- Develop techniques and tools to allow the abstraction of Java programs to allow more efficient V&V.
- Experiment with distributed and parallel search algorithms to extend the capabilities of the GVE.
- Integrate static analysis techniques with the search algorithms within the GVE to allow larger systems to be validated.
- Continue research on the development of new techniques to reduce the cost and improve the behavioral coverage of the testing required by the FAA for the certification of flight software.
- Continue research on the V&V of DEOS and its application programs in conjunction with Honeywell Research Center.
- Advanced research and development on software synthesis tools using deductive synthesis and other techniques from formal methods with the long-term goal being the synthesis of efficient, numerically stable, and customized data analysis software from a Bayesian network describing the probabilistic structure of the data analysis problem.
- Develop techniques for efficient deductive retrieval of software components and investigate its integration into a program synthesis environment.

II.A.ii.h   Problems Encountered:

None

II.A.ii.i   Problem Resolution:

None needed.

II.A.iii   <u>Advanced Planning and Scheduling</u>

II.A.iii.a   Task Summary:

The goal of this task is to extend and improve autonomous planning and scheduling technology in support of NASA missions operations, for both on-board systems and on-ground operations. This will be accomplished by extending the capabilities of the Remote Agent planning techniques, which were flight validated on board the Deep Space One spacecraft in May 1999, and by developing new solutions to problems in planning and scheduling for mission operations.

II.A.iii.b   RIACS Staff:

Ari K. Jonsson
Robert Morris

II.A.iii.c   Overview:

Traditionally, mission operators have operated spacecraft and rovers, that build and test command sequences and then upload the result to be executed on board the spacecraft or rover in question. The drawbacks of this approach are the requirement for manpower to build and test sequences, the need for deep space communications to upload sequences at regular intervals, and the lack of flexibility during execution, enforced by the predetermined sequences. As more and more missions are active at the same time, each with limited resources and manpower, these drawbacks become significant obstacles.

In recent years, autonomous techniques have been suggested and developed as possible solutions to the above-mentioned problems. The best-known result of this work is the Remote Agent, which is a complete, autonomous flight control system for spacecraft. In a unique experiment, the Remote Agent was tested on board the Deep Space One spacecraft, and performed well, despite some setbacks. One component of the Remote Agent system was an autonomous planning and scheduling engine that accepted high-level mission goals as inputs, and generated detailed operations plans that would achieve the mission goals without violated spacecraft constraints or overuse limited resources. This capability, to automatically find plans that achieve higher-level goals, is fundamental in reducing operator workload, reducing the need for deep space communications, and improving the flexibility of spacecraft and rovers during their mission tasks.

The high-level goal of this task is to continue the development of autonomous planning and scheduling techniques, to support NASA mission operations. This includes both autonomous operations on board spacecraft and rovers, and on ground autonomous plan generation as part of mission operations. The approach used to achieve this goal, is based on using constraint-based interval planning as the core plan representation and reasoning mechanism. In this approach, activities are represented as temporal intervals that are connected by constraints into a network that defines a set of plans. The Remote Agent planning utilizes this approach. However, it also extends it in important ways. For one, the Remote Agent planner arranges activities on timelines that represent the development of various spacecraft or rover components, which has the advantage of being natural to model while also simplifying one of the key activity interaction problems that appears in traditional partial-order planning methods. Another important difference is that the RA planner allows for much richer domain modeling, which in turn makes it possible to reason about such complex systems as spacecraft and rovers. In addition to building on the success of the Remote Agent planner, this task also builds on the use of constraint reasoning techniques for planning and scheduling applications in general. Research into constraint-based methods has yielded a number of important results and real-world applications in recent years. The use of such techniques can improve the efficiency, flexibility and reliability of autonomous planning and scheduling.

II.A.iii.d   Project Description:

The main component of this task is the project on constraint-based planning. Within this project, there are roughly four main research thrusts, each of which addresses an important issue in planning with constrained interval tasks in complex domains. The first is the development and implementation of a Next Generation Remote Agent planning framework. The resulting system will serve both as a research tool and as a core for future applications of planning and scheduling. The second thrust is the development of advanced constraint reasoning methods, for use in constraint-based interval planning applications. There are three main components to this goal: the development of a general framework for handling constraints and specialized reasoning methods effectively; the development of more effective techniques for dealing with specific types of constraints, such as flexible resource constraints; and research into more effective general constraint reasoning techniques, aimed at determining consequences of choices as quickly as possible. The third main research thrust is aimed at handling issues arising in controlling the search for suitable plans. There are two aspects to this control: limiting the number of options to choose from; and selecting among remaining options. To reduce the number of options, the goal is to develop and implement techniques for analyzing the planning domain to find options that are not feasible, and for finding equivalent options so that if one fails, the other can be eliminated. To better choose among the remaining options, the goal is to develop a set of techniques for specifying and deriving heuristic information. The candidate approaches for achieving this goal include the use of hierarchical task networks and the use of automated learning of heuristics from sampling. Finally, the fourth thrust of this work aims at extending the Remote Agent planning framework to handle more flexibility and uncertainty, i.e., by developing techniques to represent and reason about flexible aspects of a domain, while guaranteeing continued execution regardless of outcome. The results will be plans that are more likely to be executed successfully.

Another component of this task is work on extending the constraint-based planning paradigm to handle contingencies. This requires the constraint-based interval-planning paradigm to be extended to contingent branches in the plans. It also requires the development of techniques for executing such plans and determining if plan candidates can be successfully executed.

Finally, a third component of this task is the development of advanced constraint-based techniques for solving planning and scheduling problems. This exploration will be driven by the problems that arise from NASA applications, including constraint-based planning and scheduling, as described here above.

II.A.iii.e   Accomplishments during FY99:

The key accomplishment for fiscal year 1999 has been the progress made in designing and implementing the Next Generation Remote Agent Planner. This work is a foundation on which future research in constraint-based planning can build, while also providing an effective core planning framework that can be used in future applications of autonomous planning. Among the achievements are the following:

- A new planning framework, based on the successful Remote Agent planner, has been defined. The resulting framework is significantly simpler, both conceptually and in terms of the types of operations that are required to develop a plan. At the same time, the

framework retains all of the key generality of the original Remote Agent, and adds some generality, e.g., in the modeling of domain constraints.

- Significant progress has been made on an implementation that is based on the simpler planning framework, has been designed. The implementation design is complete and most of the modules have been implemented and tested. It is also worth noting that any of the components have been designed to be used in other reasoning applications. The most likely candidate, the constraint-reasoning component, has already been selected as a candidate for performing the constraint reasoning in another project.

- The constraint reasoning module has been defined implemented and tested on different types of constraint networks, including the networks that arise in the Deep Space One scenario. The results are promising, even though these tests have not taken advantage of some of the added generality of the new constraint-reasoning module.

- The top-level module of the new planner, which handles and reasons about the activities and constraints in a plan, has been designed, implemented and tested. As a result, the new system is capable of solving planning problems, even though some of the external modules, such as input/output of models and plans, have not been completed yet. Work is underway to test the module with a simple search engine on a realistic spacecraft science observation-planning problem.

In addition to the above-mentioned work, progress has been made in other interesting areas of constraint-based autonomous scheduling. In particular, two different types of problems were analyzed as constraint reasoning problems:

- Satellite communications scheduling involves a rapidly changing set of calls and connections, where calls are handled by assigning them valid connection paths. This is then complicated by different calls having different priority. During this year, this task provided a formulation of prioritized satellite call scheduling as a dynamic constraint optimization problem. This then led to the recognition that an existing call assignment technique was in fact a hill-climbing method. Analysis of this method then showed that its complexity was at worst linear in the number of satellite call requests in the system, thus providing performance guarantees for the method. Furthermore, experimental results showed that the hill-climbing method provided solutions that were at most 5% worse than optimal solutions.

- Cyclic scheduling problems are an interesting class of scheduling problems that involves finding schedules that can be repeated at regular intervals. This year, this task resulted in the development and testing of a new formulation of cyclic scheduling as constraint satisfaction. This result provides a general mechanism for solving cyclic scheduling problems as constraint satisfaction problems, which in turn permits the application of various different techniques for solving such problems. In addition, this makes it possible to solve cyclic scheduling problems that are embedded in other constraint problems. The formulation was tested by translating a set of benchmark problems into constraint

satisfaction problems, and then, solve them. The solver used was a simple branch-and-bound method that utilized procedural reasoning to perform certain reasoning tasks, such as keeping track of earliest possible start times and cycles, more effectively. The results demonstrated that even fairly simple techniques could solve this formulation of cyclic scheduling as constraint satisfaction.

II.A.iii.f Publications and Presentations during FY99:

1. Ari K. Jonsson, Paul H. Morris, Nicola Muscettola, Kanna Rajan, "Next Generation Remote Agent Planner". *In the Proceedings of the 5th International Symposium on Artificial Intelligence, Robotics and Automation in Space (iSAIRAS99), ESTEC*, Noordwiijk, Netherlands, 1999.

2. Denise Draper, Ari K. Jonsson, David Clements, David Joslin, "Cyclic Scheduling". *In the Proceedings of the 16th International Joint Conference on Artificial Intelligence (IJCAI99)*, Stockholm, Sweden, 1999.

3. Ari K. Jonsson, Jeremy D. Frank, "A Framework for Dynamic Constraint Reasoning using Procedural Constraints". *Workshop on Constraints in Control, part of the International Conference on Principle and Practice of Constraint Programming (CP99)*, Alexandria, Virginia, 1999.

4. Christian Plaunt, Ari K. Jonsson, Jeremy Frank, "Run-Time Satellite Tele-Communications Scheduling as Dynamic Constraint Satisfaction". *In the Proceedings of the IEEE Aerospace Conference*, 1999, Snowmass, Colorado, 1999.

5. Christian Plaunt, Ari K. Jonsson, Jeremy Frank, "Satellite Tele-Communications Scheduling as Dynamic Constraint Satisfaction". *In the Proceedings of the 5th International Symposium on Artificial Intelligence, Robotics and Automation in Space (iSAIRAS99)*, ESTEC, Noordwiijk, Netherlands, 1999.

6. David E. Smith, Jeremy D. Frank, Ari K. Jonsson, "Between Planning and Scheduling". Invited paper to a special issue of *Knowledge Engineering Review*, to be published in early 2000.

II.A.iii.g Future Plans:

The long term goals of this work are to develop autonomous technology for planning and scheduling tasks in complex domains, and to make this technology useful and available to NASA missions. In the short term, for the fiscal year 2000, the goal is to reach at least the following milestones:

- Extend the Remote Agent plan representation to increase flexibility and effectively handle uncontrollable events.
- Develop and test efficient reasoning methods for handling flexible resource usage in the constraint-based interval-planning framework.
- Develop and implement a process for reducing the plan search space, based on abstraction.

II.A.iii.h   Problems Encountered:

None.

II.A.iii.i   Problem Resolution:

None needed.

II.A.iv   <u>Bayesian Inference and Image Analysis</u>

II.A.iv.a   Task Summary:

NASA's EOS program, as well as other data collecting government agencies (such as NOAA, USGS, NIMA, etc.) and industry generate vast sets of observational data of the earth. This data is a basic resource, that can help answer such basic scientific questions as global warming, changing land use, ocean circulation and so on. In addition this data provides the raw material for answering broad policy questions, such as crop assessment, forestry planning, urban planning, etc. Also, individual users may find this data helpful in providing information about their particular farm or community. Given this range of users of remotely sensed data (and possibly in situ data as well), this raises the question how can NASA best meet the information needs of these users?

This problem of integrating information from different sensors to answer particular questions is a familiar one in remote sensing and other fields, and often is referred to as "data fusion". "Data fusion", as sometimes envisioned, is a fundamental misconception. Data should **never** be tampered with, let alone "fused". Data is what was observed, and as such cannot be changed after the event. What basic theory says is possible is to construct geophysical models of the surface and atmosphere that "best" predicts the observed data. For example, a geophysical ground model would include a point-by-point description of the topography, ground cover types and spectral characteristics, soil types, moisture content, etc. From such a geophysical model, it is possible to project what a particular patch of ground would look like from a particular viewing angle, under particular lighting conditions, with a particular camera etc. That is, given the model, one can probabilistically predict what would be observed (the expected data). The difference between what is actually observed (the data) and the expected data can be used to update the model. Thus the model acts as the central repository of all the real information in the data and is constructed from our prior knowledge of geophysical processes and how radiation interacts with the geophysical system. This central model could loosely be described as the result of "fusing" the data, but it is not itself "data" or a "data product". The Bayesian probabilistic estimation approach not only allows estimation of the most probable model given prior domain knowledge and the data, but it also estimates the uncertainty associated with the model. In particular, if this model uncertainty is high, it means that there is insufficient data/prior knowledge to pin down which model actually applies.

The Bayesian model-based data integration, outlined above, in principle solves the problem of how to integrate information form multiple sensors. However, there are many practical problems in constructing computationally efficient models, especially in view of the huge amounts of data involved. In the long run, it should be possible to integrate all satellite (or other) data into one

global model as the data is received. This updated model can then be used to project particular information required by particular users. It is also very useful for spotting changes, because these are where there is a larger than expected difference between what the model expects and what is observed. The model can also be used to compress the data, since only the differences between expected and observed need to be stored. In the short term, this technology potential can be tested in much more limited demonstrations of data integration, such as integrating the images from the same satellite on different passes. This limited case is what the Bayesian Inference group is investigating.

II.A.iv.b   RIACS Staff:

Peter Cheeseman
Robin Morris

II.A.iv.c   Technical Approach:

The basic theory behind our approach is that of inverse graphics. That is, if we knew what the ground is like, the lighting conditions, and the camera orientation and characteristics, etc., then we could predict what the camera would see (an image). This is the standard computer graphics problem. However, we have the inverse problem--we know what the images are, and we want to find the most probable ground truth (surface) that would have generated them, assuming we know the lighting conditions and camera characteristics. The most important (and difficult) part of this process is recovering the camera orientation and position for each image. To do this, we must register all the images with respect to each other to an accuracy of a small fraction of a pixel; this registration tells us how an image maps onto the ground truth model we are building.

Our initial ground model is formed by letting each pixel "vote" on what the corresponding ground position should be depending on how much that ground position contributed to that pixel. This initial ground model is then used to project what each image should be (i.e., predict each pixel value). The differences between the predicted pixel value and the observed value are used to update the ground model until it cannot be further improved. This procedure produces an increase in both spatial resolution and gray-scale resolution (super-resolution). We are now able to generate images from a high-resolution ground model, and then reconstruct the surface from these images at a resolution much higher than the input images. We are currently trying to extend the system so that it can learn the camera parameters and lighting conditions while simultaneously inferring the unknown surface.

II.A.iv.d   Accomplishments during FY99:

We have been successful in reconstructing an unknown surface from low-resolution images of that surface, even though the images are taken from different viewpoints and under different lighting conditions. This is the first time this has been achieved to our knowledge. The results of this super-resolution can be seen on our web page:
http://ic-www.arc.nasa.gov/ic/projects/bayes-group/super-res/2d/mpf

II.A.iv.e   Publications and Presentations during FY99:

"A Bayesian Approach to High Resolution 3D Surface Reconstruction from Multiple Images", by R. Morris, P. Cheeseman, V. Smelyanskiy and D. Maluf.

II.A.iv.f   Problems Encountered:

None.

II.A.iv.g   Problem Resolution:

None needed.

II.A.iv.h   Visiting Scientist - Bayesian Inference for 3D Superresolution

Anil Kokoram visited RIACS and worked with this task group. This is his report.

II.A.iv.h.1   Task Summary:

In order to infer the depth map of a surface from a series of multiple views, a prior model for the surface is advantageous. It is known that multidimensional Autoregressive (AR) models provide an analytical basis for high fidelity image reconstruction in general. The Task was to investigate the use of an AR process as a model of the local smoothness of a surface.

II.A.iv.h.2   RIACS Staff:

Anil C. Kokaram

II.A.iv.h.3   Overview:

The Superresolution group has been developing algorithms which combine the information from multiple views of a scene (a planet surface) to generate a higher resolution composite image than any of the individual observed images.

Recently, the work has been extended to recover a depth map simultaneously at this higher resolution. The key to the process is to design a generative model for the observed views. This implies defining the surface in 3D space at a high resolution and synthesizing the images of each view, given the camera parameters. The task then is to estimate that underlying surface model which can generate the actual images observed at the different camera viewpoints. The whole analysis has been articulated under a Bayesian framework.

There are several issues involved in the design of the model and the solution of the unknowns. The synthesis process is essentially a computer graphics problem, and the group employs a triangular mesh for defining the surface heights along with assumptions of Lambertian reflectance etc. The estimation problem is difficult because of the large number of variables and the highly non-linear nature of the graphics synthesis process. Markov Chain Monte Carlo algorithms could be used, but the group has had considerable success with deterministic search strategies which involve multigrid estimation and linearization of the system about some starting estimate for the surface heights. A key issue in accurate estimation is the design of good priors for the local map smoothness and, in effect a good prior for the surface structure. Surprisingly good results have

been generated by using a prior which encourages maximal smoothness. Nevertheless, this model is not necessarily realistic, especially given the stochastic nature of a depth map at certain scales.

Dr. A. Kokaram has extensive experience in the area of image and video reconstruction for the purpose of missing data reconstruction. He has employed linear stochastic models (Two dimensional autoregressive (2D AR) processes) as a prior on local image structure. The model controls the correlation of image intensities at a local scale. This model was used for reconstructing image data in missing gaps. This implies that the AR process is a good model for such structures. For the 3D superresolutuion problem, it is clear that the better the prior surface model that is used, the better an inference for depth can be made.

Given the relatively straightforward statement of this model, it is logical to consider depth values as intensities in the 2D AR model (as it is normally used), and so investigate a simple extension of the current smoothness prior used to the case of this `stochastic' smoothness prior.

Therefore, the intended contribution (during August 1999) was to consider the 2D AR model in this context and assess its use within the estimation framework for 3D superresolution.

II.A.iv.h.4   Project Description:
Assess the accuracy with which the AR process could reconstruct missing data in the depth map. Incorporate the process as a smoothness prior in its simplest form.

II.A.iv.h.5   Accomplishments during FY99:
The process of interpolation of missing data when the data is a set of depth samples was successfully implemented. The results indicate that the AR model is a viable model for the local structure in such depth maps as were being used by the group at the time. However, the scale of the underlying data dictates the order of the model which is most appropriate, thus adaptive model order selection (a problem not considered) becomes important in this context.

An adaptive smoothness process was implemented which in effect, measured the correlation between depth samples and relaxed the smoothness constraint accordingly (similar to a first order AR model). There was not enough time for exhaustive testing on site.

II.A.iv.h.6   Publications and Presentations during FY99:
Dr. Kokaram gave a presentation of his video reconstruction work at a RIACS seminar during August 1999.

II.A.iv.h.7   Future Plans:
Tests with the AR model have emphasized the importance of a knowledge of "scale" in the estimation of these depth maps. Future work dealing with the specification of a better prior for the surface will concentrate on this issue by designing models which incorporate a notion of scale and local directionality. This would allow large features on the surface to be estimated quickly and finer surface details to be reconstructed depending on the need of the user.

II.A.iv.h.8   Problems Encountered:

There was a problem with the timing of the visit, as it coincided with an important graphics conference (SIGRAPH '99), and a surprising call for interview in the U.K. This meant that there was less contact with personnel available in the initial week of the visit. This problem could not have been foreseen, however.

II.A.iv.h.9   Problem Resolution:

Dr. Morris did his best to prepare Dr. Kokaram for the visit by way of introductory materials and background information available on his arrival. As a result, the visit was immensely successful and beneficial.

II.A.v   Advanced Onboard Autonomy

II.A.v.a   Task Summary:

Current and future planetary exploration missions require robotic spacecraft that are capable of performing complex tasks with minimal communication with ground-based operators. The research conducted under this task supports the development of advanced onboard autonomous control architectures and methodologies, which enable future NASA planetary exploration missions.

II.A.v.b   RIACS Staff:

John Loch

II.A.v.c   Overview:

Future exploration missions require new methods of commanding and controlling robotic agents on distant planetary and asteroid surfaces. Communication with robotic agents operating on distant planetary surfaces is constrained by bandwidth, opportunity, and time delay. Bandwidth and opportunity constraints can be improved by building additional communication infrastructure, but the time delay imposed by the limited speed of light is here to stay.

What this means for the operation of planetary surface rovers is that at a minimum they must be capable of autonomously executing human designated command sequences on the order of tens of minutes between command cycles for operation on Mars and nearby asteroids. Rover operations on the moons of Jupiter, such as Europa, will require robust execution of command sequences on the order of hours between command cycles.

The objective of the research conducted under this task is to develop robust, fault-tolerant autonomous control architectures capable of executing human designated command sequences. The research conducted under this task is being applied to two NASA projects: the Personal Satellite Assistant (PSA) project and the Robust Task Execution project.

The Personal Satellite Assistant (PSA) project is a NASA project to develop a free-flying robot that will serve as a virtual astronaut assistant onboard the International Space Station. The PSA robot will be used to monitor environmental conditions, assist astronauts in performing

experiments and serve as a telepresence device for ground-based operators via an onboard videoconferencing system.

The Robust Task Execution project is a recently funded NASA activity that is exploring the use of reinforcement learning methods to improve the performance of autonomous rover navigation in natural terrain. The successful application of the research conducted under this task can lead to enhanced autonomous navigation algorithms, which will enable long-range Mars rover operations, planned for the 2007 and 2009 Mars rover missions.

II.A.v.d   Project Description:
- Development of a reinforcement learning architecture for robust navigation of mobile robots in both 3-dimensional spaces (PSA) and natural terrain.
- Research methods for handling large state space exploration for learning optimal navigation strategies.
- Research methods for automatic discretization of robotic sensors for optimal performance.
- Research methods for incremental learning of tasks.
- Research methods for transferring learned strategies between tasks.

II.A.v.e   Accomplishments during FY99:
- Formulated a research and development plan for the testing and validation of the PSA robotic assistant concept from initial design to flight testing onboard the space station.
- Completed the design of the PSA robot prototype hardware and software.
- Designed and implemented a custom embedded computer architecture for the autonomous operation of the PSA robot.
- Collaborated with Prof. Howie Choset from CMU to develop the onboard autonomous path planning and control software for the PSA robot.
- Authored two proposals for funding for NASA Ames Research Center's Autonomy and Robotics Area.
- Initiated a research discussion group for discussing machine learning research for space exploration applications.

II.A.v.f   Publications and Presentations during FY99:
None.

II.A.v.g   Future Plans:
- Development of a rover-terrain simulator.
- Development and testing of reinforcement learning architecture for autonomous mobile robot navigation using rover-terrain simulator.
- Design and construction of microgravity PSA testbed prototype.
- Development and testing of autonomous 3-D navigation algorithms.
- Development of 3-D position and orientation state estimator using stereo vision and inertial sensors.

II.A.v.h   Problems Encountered:

None.

II.A.v.i   Problem Resolution:

None needed.

II.A.vi   <u>Fusion and Compression Focus Image Series</u>

II.A.vi.a   Task Summary:

Two tasks were accomplished during 1999. The first task involved segmentation and analysis of echocardiograms and ultrasound images using stochastic relaxation and was conducted at Center for Bioinfomatics at NASA Ames and is discussed at section below. The second task was 3D reconstruction, fusion and compression of image series from the Robotics Arm Camera (RAC), which will be flown on missions 2001 and 2003.

II.A.vi.b   RIACS Staff:

Esfandiar Bandari

II.A.vi.c   Overview:

Microscopic imagers, like the Robotic Arm Camera (RAC) and MECCA are being planned for the Mars Surveyor '01 and '03 missions. A common problem with these imagers is that they have a limited depth of field in near regions causing out of focus areas within each image. As a result several images are captured which cover different depths of field in the imagery with the remainder of the image being out of focus. Given the constrained bandwidth associated with these missions, as well as a need for clear visualizing of the complete scene from the returned images, our goal was to:

1. Significantly reduce the data volume by autonomously identifying and returning only the focused portions of each image. This is particularly important in conjunction with the upcoming JPEG 2000 specification, which allows variable compression ratios within one image.
2. Create a single in focus image for analysis and visualization.
3. Create a 3D reconstructed model of the scene.

We were able to successfully complete these tasks, presenting our results in the two publications listed below. Three-dimensional reconstruction from zoom is an interesting area of research with applications to landing safety analysis from descent imagery and one, which we are planning to pursue in upcoming proposals.

II.A.vi.d   Project Description:

The key issue is to find a good metric for autonomously distinguishing in focus regions from blurred portions of an image. Other problems facing this project were variable spatial blur within individual image blurs and spatial scaling between every two frames. We examined various algorithms, including statistical methods, filtering techniques and those based on maximum

entropy. The novel focus operator developed by the author is based on a spatial high frequency detector. The output of this operator is an image with regions of high frequency (low blur) with greater magnitude than blurred regions. Using statistical and scale space analysis we are then able to identify regions of each image that are most in focus.
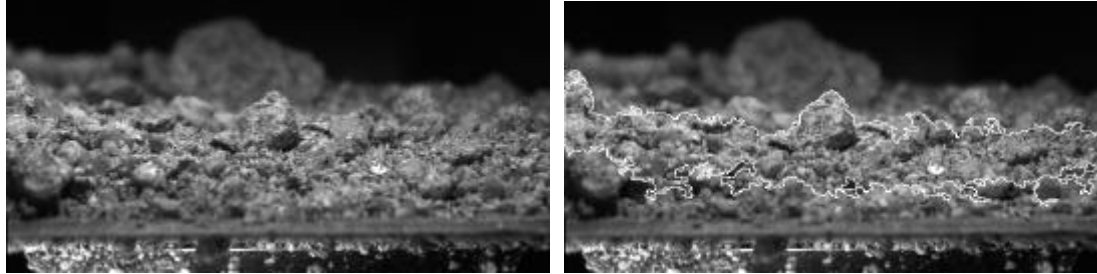
The left figure shows a sample image from the focus image series captured by RAC. The right image shows the boundary region that is marked as in focus. Using this approach we were able to achieve an 82% compression ratio in 7 images captured by RAC.

II.A.vi.e   Accomplishments during FY99:

This researched involved several novel algorithms including a new spatial blur detector operator. We were able to autonomously segment the images and achieve an 82% spatial compression from blur elimination process alone. As proof of concept we reconstructed a single in focus image from the 7 images originally captured by RAC. The resulting image presented a complete and clear representation of the scene.

II.A.vi.f   Publications and Presentations during FY99:

1. V.C. Gulick, R.L. Morris, E. Bandari, M.A. Ruzon, T.L. Roush, "Autonomous Image Analysis for Future Mars Missions", Presented at the *American Geophysical Union Meeting*, San Francisco, CA, December 1999.

2. V.C. Gulick, R.L. Morris, M.A. Ruzon, E. Bandari, T.L. Roush, "Autonomous Onboard Science Image Analysis for Future Mars Rover Missions", *31st Annual Meeting of the Division of Planetary Science*, Padova, Italy, October 1999

II.A.vi.g   Future Plans:

- Participate in the field test of the algorithms.
- Secure funding for 3D structure from zoom and 3D structure from motion to continue this work.

II.A.vi.h   Problems Encountered:

None.

II.A.vi.i   Problem Resolution:

None needed.

II.B   Human-Centered Computing

II.B.i   Workspace Analysis

II.B.i.a   Task Summary:

The Work Practice Analysis Task is part of the Human-Centered Computing focus area of Ames Research Center (COE-IT). The focus of the research is on understanding how people and systems are interconnected in practice. To accomplish this, research is conducting in the areas of work systems analysis and evaluation and the development of computational modeling tools for simulating how people collaborate, communicate, and work within their environment.

II.B.i.b   RIACS Staff:

Maarten Sierhuis
John O'Neill

II.B.i.c   Overview:

The research conducted under this task is an integral part of the Work Systems Analysis and Evaluation group in Code-IC. The group's research philosophy is based on the view that Human-Centered Computing is a software engineering methodology. This methodology is based on the scientific study of cognition in people and machines, especially understanding the differences between perceptual-motor/cognitive/social aspects of people and present-day computer systems with the objective of developing computer systems that fit human capabilities and practices by exploiting and improving AI programming methods.

II.B.i.d   Project Description:

Three projects were conducted under this task during FY99:

- Brahms

The objective of the Brahms project is to research how we can model the work practices of people and systems as an integral part of a human-centered systems engineering approach. Brahms is a multi-agent modeling and simulation environment for simulating the situated work processes and activities of organizations, including people and systems.

- Logging/Handover Tools for Mission Control

JSC is developing "ops assistants" to support the work of flight controllers in the International Space Station's Mission Control Center. The first "ops assistant" to be developed is a tool to support the logging and handover practices of flight controllers. This project is a collaboration between the Mission Operations Directorate at JSC, the Engineering Division at JSC, the Astronaut Office at JSC, the Work Systems Design and Evaluation Group at NASA Ames, and Ohio State University.

- Studying Delays at United Airlines

A major objective of NASA's Aviation Program is to improve the efficiency of the U.S. Airspace system. A major contributor to the efficiency of the U.S. airspace system are the airlines,

, TM99-0008 (Brahms Project Document), Version 2.0, December 1999.

3. Sierhuis, M. (1999). "Modeling and Simulation of the Apollo 12 ALSEP Offload", Invited Talk at the *Institute for Research on Learning*, Menlo Park, December 1999.

4. Sierhuis, M. (1999). "Modeling and Simulation of the Apollo 12 ALSEP Offload: Brahms Modeling and Simulation Experiment Results", (not yet published).

5. Selvin, A.M., M. Sierhuis. (1999). "Argumentation in Different CSCA Project Types", *Proceedings of the Computer-Supported Collaborative Argumentation for Learning Communities Workshop, Computer Supported Collaborative learning (CSCL) Conference* '99, Stanford University, CA. http://kmi.open.ac.uk/sbs/csca/cscl99/papers.html.

6. O'Neill, J. & Wales, R. "CSCA Issues Raised by Mission Control for the International Space Station", *Proceedings of the Computer-Supported Collaborative Argumentation for Learning Communities Workshop*, 1999 http://kmi.open.ac.uk/sbs/csca/cscl99/papers.html.

7. O'Neill, J. "Implications of Learning by Peripheral Observation for Virtual Environments", *Proceedings of the Third International Conference on Cognitive Technology*, 1999.

8. O'Brien, F. & O'Neill, J. "Decision-Making and Currency for C2 Systems", *Proceedings of the Fifth International Command and Control Research and Technology Symposium*, 1999.

II.B.i.g   Future Plans:

- Brahms
    1. Implement and Test the Brahms Java-based Virtual Machine.
    2. Develop a user-friendly Brahms model-builders environment.
    3. Develop a Java-language interface to Brahms.
    4. Apply Brahms as a multi-agent planning and execution system for intelligent agent systems.
    5. Extend Brahms with Virtual Reality Display capabilities for end-user interaction.
    6. Apply Brahms as a knowledge-based requirements engineering tool for human-centered design.

- Studying Delays at United Airlines
    1. Create models of United Airline's ground operations, identifying communication and information flow "gaps" and the effects of rework caused by delays.
    2. Develop new technologies to address the communication "gaps" between players in delay situations, creating new sources of information for reasoning about delays.
    3. Integrate our models and findings with work being conducted in the SMS and Technical Area and Future Flight Central facility to aid designing a more efficient airspace system.

- Studying Rework at the International Space Station Mission Control Center
  JSC's Mission Control Operations have traditionally been focused on short-duration operations, with all members of the flight control team centrally located in one facility. Long-duration missions and a flight control team distributed across many locations in different time zones around the world will characterize Mission Control operations for the International Space Station. We are proposing a new study to investigate the computer support required for the International Space Station Mission Control Center, particularly in breakdown and rework situations.

II.B.i.h   Problems Encountered:

1. John O'Neill has not been permanently badged by NASA/Ames this year due to Export Control Office problems.
2. The Brahms project lost a developer in March '99. We were not able to identify appropriate candidates for hiring. When we identified appropriate candidates three months later, we were not able to hire a new developer due to loss of funding in the meantime. Because of the loss of the resource the development of the Brahms Virtual Reality Environment was stopped.

II.B.i.i   Problem Resolution:

1. No resolution as of yet.
2. Submitting new proposal(s) for additional funding.

II.B.ii   RIACS Language, Interfaces and Speech Technology

II.B.ii.a   Task Summary:

The group's primary focus is on Spoken Dialogue Natural Language Interfaces (SDNLIs) to semi-autonomous agents. Work builds on top of established speech and language technology, in particular, the Nuance recognizer and the SRI Gemini and Open Agent Architecture systems. Our research over the last year has concentrated on contextual interpretation, portability, dialogue management and interface level issues.

II.B.ii.b   RIACS Staff:

Manny Rayner
Beth Ann Hockey
Frankie James

II.B.ii.c   Overview:

Speech recognition and natural language processing have been identified as new growth areas for NASA Ames. Commercially available tools for development of spoken language interfaces are for the next few years likely to remain relatively primitive, and be aimed mainly at menu-driven applications based on phrase spotting and slot-filling. The high-level objective of our work is to create a generic software platform that will allow NASA personnel to develop spoken dialogue natural language interfaces to semi-autonomous robots, advanced planning systems, and other complex agent-like software. We are developing and integrating relevant component technologies, and testing them in prototype spoken language applications.

II.B.ii.d   Project Description:

The starting points for our theoretical and practical work on spoken language interfaces have been the Nuance, Gemini, and OAA systems. Our first task has been to find out how to use and extend these packages to build applications potentially interesting to NASA.

- Contextual interpretation

The software tools named above make it relatively easy to implement context-independent aspects of spoken language processing. As spoken language interfaces become more sophisticated,

however, it becomes increasingly important to support various kinds of contextual interpretation: the meaning of what the user says to the system normally depends on the preceding dialogue. We have given high priority to the task of extending the Gemini system in this direction, and developing a methodology for building spoken language dialogue systems which treat discourse context as an integral part of the interpretation model.

- Portability

Constructing a serious speech interface is a large task, and if possible we want to be able to reuse resources. We are building on previous work in this direction carried out at SRI International under the Spoken Language Translator project. The basic idea is to center development on a single general grammar, which is semi-automatically adapted to each particular application.

We are developing a general grammar conforming to these principles, which incorporates ideas from the CLE and XTAG projects, and implementing versions of the SRI adaptation methods for use in our current projects. The adaptation methods are described in detail in the book "The Spoken Language Translator" (editors: Rayner et al), to be published next year by Cambridge University Press.

- Dialogue management and interface level issues

There has to date been a striking lack of interaction between the HCI and speech understanding communities. We are making efforts to incorporate insights from the HCI field in the design of our prototype spoken language interfaces, in particular recent research by Clifford Nass and Byron Reeves that has studied social reactions to computer technology.

- Infrastructure

RIALIST is a new group, and during this first year we have been obliged to spend a considerable proportion of our time on building up a computational infrastructure capable of supporting our research. This has included acquisition and installation of hardware and
software; learning to use necessary basic tools, and implementation of core software.

II.B.ii.e   Accomplishments during FY99:

We began our work by developing some small applications using the Nuance toolkit, partly to get a feel for the limitations of the current COTS technology, and partly to give people at NASA a sample of how this technology could be applied in domains of potential interest to them. Three such applications were developed, for the Personal Satellite Assistant, Surface Movement Advisor, and Science Desk projects. In each case, the application used a vocabulary on the order of a few dozen words, a context-free grammar with about ten to twenty rules, and a visual display; semantic processing essentially consisted of slot filling. Though simple, these initial demos succeeded very well in their aim.

We discovered that it was difficult to extend the scope of the Nuance-based systems to include more complex strategies for semantic interpretation and dialogue management. This led us to develop a second and much more sophisticated version of the Personal Satellite Assistant interface demo, implemented on top of the Gemini system. The basic approach is to transform the input speech signal through successive levels of representation corresponding roughly to linguistic

knowledge, dialogue knowledge, and domain knowledge. The final representation is an executable program in a simple scripting language equivalent to a subset of CSHELL.

At each stage of the translation process, an input is transformed into an output, producing as a byproduct a "meta-output" which describes the nature of the transformation performed. Consistent use of the output/meta-output distinction permits a simple and perspicuous treatment of apparently diverse topics including resolution of pronouns, correction of user misconceptions, and optimization of scripts. The methods are all essentially domain-independent in nature, and could readily be ported to other applications.

II.B.ii.f   Publications and Presentations during FY99:

The group is too new to have had time to produce any publications this year. Several papers are in preparation, and will be submitted to workshops and conferences to take place during FY2000.

II.B.ii.g   Future Plans:

During the coming year, we aim to continue with our basic plan of constructing a general framework for implementation of advanced spoken natural language dialogue interfaces. We are presently discussing with various NASA groups the possibility of applying our work to their domains. In particular, we hope to be able to link up with at least one of the Remote Agent, Science Desk, or BRAHMS projects. Our concrete goals in terms of developing core functionality are the following:

- We will continue with our current work on contextual interpretation of spoken dialogue. Most obviously, we need to extend the framework to include not only interpretation of overt commands, but also of a wider range of question-types. Other important tasks are to improve coverage of reference resolution (interpretation of pronouns etc.), and to implement a more detailed model of discourse state.

- We have so far only just begun work on portability aspects. During the next year, we will integrate into the Gemini environment the work carried out under the Spoken Language Translator project, and expand our central domain-independent grammar. The resulting system should vastly reduce the time required to implement a flexible high-coverage grammar for a new domain.

- The idea of using machine-learning methods to find rational ways to combine multiple sources of information is currently generating a great deal of activity in the computational linguistics community. We plan to investigate applications of this idea to topics including dialogue management and reference resolution. It is possible that some of this work will be carried out using Bayesian search methods developed within Peter Cheeseman's group.

- Speech interfaces to robots and other embedded systems raise important safety issues. If the user is expected to control the robot exclusively through the speech modality, it is imperative to have some kind of speech-based "emergency switch", which can be used to stop a misbehaving system.

- It seems improbable to us that conventional word-based recognition technology will be able to offer sufficiently reliable performance in a near to medium-term perspective. We will investigate the idea of triggering emergency behavior not off the actual words spoken, but rather off prosodic features such as pitch, amplitude and duration. Approaches to investigating such prosodic features are discussed in "The Interpretation and Realization of Focus" (Hockey 98) and "Prosody and the Interpretation of Cue Phrases" (Hockey 92).

- Last, but not least, we will carry out controlled user studies to evaluate the above and other related work. The empirical methodology for these studies will in part build on that developed in Frankie James' 1998 dissertation, "Representing Structured Information in Audio Interfaces".

II.B.ii.h   Problems Encountered:

The main problems we have encountered have been those one would expect in a new group working in a new domain; we need to find a niche in the local ecology. Apart from that, the only item that stands out is that we are uncomfortably dependent on SRI research software (in particular the Gemini system), not all of which is necessarily well documented or easy to understand.

II.B.ii.i   Problem Resolution:

People at NASA have been very supportive in helping us find our role here, and we feel we are making excellent progress toward doing so. With regard to the SRI software, it appears that we may be able to hire an ex-SRI employee who was instrumental in implementing the Gemini system. This would be an ideal way to resolve the problem.

II.C   High Performance Computing and Networking

II.C.i   NGI/NREN

II.C.i.a   Task Summary:

The NREN/NGI project is NASA's contribution to the Federal NGI program. The focus of the NREN project is to develop and demonstrate applications that will revolutionize the conduct of NASA science. RIACS supports many of the major activities within this project.

II.C.i.b   RIACS Staff:

Marjory J. Johnson

II.C.i.c   Overview:

NREN is a next-generation network testbed that peers with NGI testbeds sponsored by other Federal agencies and with the university-led Internet2 testbed to provide a nationwide network testbed for conducting network research and demonstrating revolutionary applications. RIACS provides support for the NGI/NREN program at NASA Ames, including conducting basic

networking research, participating in projects to develop new networking technology and revolutionary network applications, and assisting with performance engineering.

II.C.i.d   Project Description:

M. Johnson is part of the management team within NREN. Components of the RIACS task include assisting in the formulation of NGI/NREN projects, participating in specified research projects for the development of new networking technology, interfacing with the university research community, participating in development of revolutionary network applications, and publishing and presenting results at conferences.

II.C.i.e   Accomplishments during FY99:

- Bridging the Gap Workshop:

NREN hosted a workshop entitled Bridging the Gap from Networking Technologies to Applications on August 10-11, 1999. The objective of the workshop was to facilitate the transfer of next-generation networking technologies to leading edge revolutionary applications. Results from the workshop include summaries of the current status of Quality of Service (QoS), advanced multicast, and security technologies; roadmaps identifying expected technology developments within the next one to five years; and guidance for application developers to assist them in inserting new technologies as they become available. M. Johnson co-organized the workshop, led the multicast breakout group on the second day of the workshop, and was one of the primary authors of the final report of the workshop. This report is located at www.nren.nasa.gov/workshop4.html. A summary of the results of this workshop will be presented at the HPCC/CAS (High-Performance Computing and Communications / Computational Aerosciences) workshop in February 2000.

- Virtual Collaborative Clinic Application:

M. Johnson was part of an NREN team that worked with the NASA Center for Bioinformatics to demonstrate the Virtual Collaborative Clinic (VCC) application across wide-area networks in May 1999. Three-dimensional high-resolution images of the human body (obtained via medical-imaging devices such as MRIs, CT scans, and echocardiography) were transmitted across the network in near real-time while medical personnel at the various sites rotated and manipulated the images, collaboratively diagnosed patient conditions, and performed virtual surgery. The primary network challenge was enabling high-bandwidth multicast, up to 30 megabits per second, across multiple network domains. Subsequent to the joint May demonstration, NREN has demonstrated networking aspects of the VCC application at INET '99 in June 1999 and at the Bridging the Gap Workshop in August 1999. A demonstration is also scheduled for Supercomputing (SC) '99 in November 1999. M. Johnson was the primary author of a paper entitled "Using the NREN Testbed to Prototype a High Performance Multicast Application," which has been accepted for presentation at SC '99.

- University Collaborations:

M. Johnson coordinates activities pertaining to NREN research grants to universities; she has made several site visits to the various universities to make a presentation about NREN activities and applications, and to discuss collaboration with the university PIs. In addition, M. Johnson assisted Professor Willard Smith of Tennessee State University in preparing a proposal that

received a NASA FY99 Partnership Award for Innovative and Unique Education and Research Projects. M. Johnson organized an NREN PI day on August 12, 1999 to bring together the university PIs with NASA personnel from across the country to enable networking personnel at all NASA sites to benefit from grant activities funded by NREN. The university PIs presented summaries of their research projects during this meeting.

- Other NREN Collaborations

M. Johnson helped develop a proposal for NREN to participate in the Internet2 QBone testbed and a proposal for joint technology activities between NREN and the Earth Science Data and Information System (ESDIS) project, located at NASA Goddard Space Flight Center.

M. Johnson is a member of the Networking Research Team (NRT) within the Federal Large Scale Network group. The purpose of the NRT is to coordinate research activities among all the Federal agencies that are NGI partners. Within the NRT we are developing a virtual seminar series to promote collaboration on technology development and to publicize the research aspect of NGI activities.

- Other Activities:

M. Johnson served on the program committee for the IFIP Workshop on Protocols for High-Speed Networks '99, the European Conference on Multimedia Applications (ECMAST '99), and the technical papers committee for Supercomputing '99 Conference.

II.C.i.f   Publications and Presentations during FY99:

1. Marjory Johnson, Matthew Chew Spence, and Lawrence Chao, "Using the NREN Testbed to Prototype a High Performance Multicast Application," accepted for presentation and publication at *SC '99*.

2. Marjory Johnson, Richard desJardins, Doug Montgomery, William Johnston and Matthew Chew Spence "Bridging the Gap Workshop report", Served as animator for multicast discussion at *IFIP Sixth International Workshop on Protocols for High Speed Networks (PfHSN '99)*, August 1999.

3. Marjory Johnson, Presented talk "NREN Project Overview," at the following universities: UCLA, June 1999, University of Illinois, Urbana-Champaign, July 1999, Naval Postgraduate School, September 1999.

4. Represented "NREN" at *CENIC '99 Conference*, May 1999; member of a panel to discuss funding opportunities for network research.

5. Marjory Johnson, Presented RIACS Seminar "NREN Contribution to the Virtual Distributed Clinic," February 1999. This chalk talk was a joint presentation about the VCC application with Muriel Ross and Xander Twombly.

II.C.i.g   Future Plans:

A student intern will be added to the RIACS task in FY 2000. Networking technologies that will be addressed next year include gigabit to the desktop, QoS middleware, and measurement and monitoring of different classes of traffic flows. In addition NREN grant activities will be more closely integrated with NREN core activities.

II.C.i.h   Problems Encountered:

None

II.C.i.i   Problem Resolution:

None needed

II.C.ii   __ARCLAN 2000__

II.C.ii.a   Task Summary:

ARCLAN was established to create an end-to-end data communications service, which will accommodate the needs of Ames into the next century. The goal of the ARCLAN 2000 Project is to provide end-to-end, high-quality, fully responsive, cost-effective, and reliable communications services to satisfy ARC's comprehensive and evolving communications requirements.

II.C.ii.b   RIACS Staff:

David L. Gehrt

II.C.ii.c   Overview:

The focus of this task is on ensuring and enhancing the reliability of the Ames network. The emphasis is on network and host security, especially the security of the authoritative name servers for the arc.nasa.gov zone.

II.C.ii.d   Project Description:

- Name server operations and security as a result of the significant changes being made to the Berkeley Internet Name Domain Server software.
- Data Base administration to improve monitoring capability of threats to systems at Ames.
- Research into design and operation of firewalls using several different designs and operating systems.

II.C.ii.e   Accomplishments during FY1999:

- Name Server Support

There are a total of four authoritative name servers serving Ames, three of which are administered by RIACS personnel. These name servers are critical to the use of the network, a point made clear when a brief network outage deprived user systems of access to the name servers and hence to other systems on ARCLAN and on the Internet.

As an example of the steps taken to improve host security on the name servers, access of the systems is now restricted to Secure Shell (SSH) for logins, improved access control to other required daemons, including Remote Procedure Call daemons. Access available to persons has also been restricted on the basis of those persons who have an actual need to have access to the systems.

This year all attempts at unauthorized access to the name servers have been logged and retained in a relational database not located on any of the authoritative name servers. This database is routinely scanned for overtly hostile threats, and steps are taken to abort future similar attempts.

- Improved Network Security

It has become clear that the wide availability of network access has increased the threats to computer systems attached to the Internet including those at Ames. One of the principal ways crackers obtain lists of systems to attack is by obtaining the zone files from an authoritative name server for the zone. That is not possible from the authoritative name servers for the arc.nasa.gov zones.

The logs referred to above generated on the name servers also have served as an early warning system for threats against other systems on the ARCLAN network.

The placement of a firewall between ARCLAN and the Internet has created a number of problems and led to the investigation of the various aspects of firewalling using a single Linux system running as a firewall. It has become clear that the problems being experienced with the Ames Firewall are not endemic to firewalls, but are the result of a poor design of the related hardware or in the configuration of the software on the firewall systems. As this FY comes to an end the problems with the Ames firewall remain unresolved and as a side effect the firewall is not totally effective.

- Ames Research Center center-wide Email

During the early part of this FY I had become involved with trying to find a solution to the serious problems being experienced by those people who were receiving e-mail on the Center's Mail hub machine. Over the course of less than two weeks, working with a senior staff member of the then support service contractor I was able to diagnose and cure the problems the users were experiencing.

II.C.ii.f   Publications and Presentations during FY1999:
None

II.C.ii.g   Future Plans:
In addition to general support of ARCLAN, likely focus will be on:
- Secure DNS
- PKI (Public Key infrastructure)

II.C.ii.h   Problems Encountered:

None.

II.C.ii.i   Problem Resolution:

N/A

II.C.iii   Device Modeling

Visiting scientists collaborated with NASA personnel in several areas of applications of high performance computing to device modeling.

II.C.iii.a   Device Process Modeling in Microelectronics

II.C.iii.a.1   Task Summary:

The purpose of this research is to develop novel concepts in nanotechnology for NASA's future needs on electronics, computing, sensors, and advanced miniaturization of all systems; and to develop highly integrated and intelligent simulation environment that facilitates the rapid development and validation of future generation electronic devices as well as associated materials and processes through virtual prototyping at multiple levels of fidelity.

II.C.iii.a.2   RIACS Staff:

Sankaran Venkateswaran (Visiting Scientist)
Shyam Ramalingam (Visiting Scientist)
Tae Oh Young (Visiting Scientist)
Dmitrious Maroudas (Visiting Scientist)

II.C.iii.a.3   Overview:

This research assists in fulfilling NASA mission needs related to:
- Onboard computing systems for future autonomous intelligent vehicles;
- High performance computing (Tera- and Peta-flops);
- Revolutionary computing technologies;
- Smart, compact sensors, ultrasmall probes;
- Advanced miniaturization of all systems;
- Revolutionary computing technologies;
- Smart, compact sensors, ultrasmall probes;
- Advanced miniaturization of all systems;
- Microspacecraft;
- Thinking' spacecraft, and;
- Micro-, nano-rovers for planetary exploration.

II.C.iii.a.4   Project Description:

RIACS activity in this area for FY99 consisted of:

- Computer Modeling of Deposition of SI Films using PECVD;

- Preconditioning Methods for Application to Semi-conductor Materials using Neural Networks;
- Modeling of Carrier Transport Effects on High Frequency Devices.

II.C.iii.a.5   Accomplishments during 1999:

Within the first of the three activity areas, a computational analysis was of growth of a-Si:H films on H-terminated Si (001)-(2(1) substrate in the temperature range 500 K (T (773 K through molecular-dynamics simulations of repeated silyl-radical impingement. The key reactions during film growth were identified through visualization of atomic trajectories and their energetics obtained through energy minimization techniques coupled with ab initio calculations based on density functional theory. Reaction mechanisms were identified leading to abstraction of hydrogen atoms from the surface through Eley-Rideal and desorption mechanisms, adsorption of radicals onto the deposition surface accompanied by breaking of surface Si-Si bonds, migration of adsorbed surface species, H transfer between surface Si atoms, and formation of clusters such as disilane through reaction of adsorbed surface species and their subsequent desorption. The energetics of the important reactions were analyzed and seen to be in agreement with ab initio calculations. The deposited films were structurally and chemically characterized and the surface hydride distribution is in good agreement with corresponding experimentally grown films. The same reactions are seen to occur over the temperature range studied, although reactions such as H abstraction and desorption of surface species occur less frequently at lower temperatures; the surface population of higher (di- and tri-) hydrides is greater at lower deposition temperatures.

Silylene (SiH_2) radicals created by electron impact dissociation of silane in reactive gas discharges can play an important role in plasma deposition of amorphous and nanocrystalline silicon thin films. We have implemented a systematic computational analysis of the interactions of SiH_2 radicals with a variety of crystalline and amorphous silicon surfaces based on atomistic simulations. The hydrogen coverage of the surface and, hence, the availability of surface dangling bonds is shown to exert the strongest influence on the radical-surface reaction mechanisms and the corresponding reaction probabilities. The SiH_2 radical reacts with unit probability on the pristine Si(001)-(21) surface which has one dangling bond per Si atom; upon reaction, the Si atom of the radical forms strong Si-Si bonds with either one or two surface Si atoms. On the H-terminated Si(001)-(21) surface, the radical is found to react with a probability of approximately 50%. The SiH_2 radical attaches to the surface either by forming two bonds with Si atoms of adjacent dimers in the same dimer row or through Si-Si bonds with one or both atoms of a surface dimer. In addition, the SiH_2 radical can attach in the trough between dimer rows, forming two Si-Si bonds with second-layer Si atoms. The energetics and dynamics of these surface reactions are analyzed in detail. A reaction probability of approximately 70% is calculated for SiH_2 radicals impinging on surfaces of hydrogenated amorphous silicon (a-Si:H) films with varying concentrations of hydrogen. Recent experimental measurements have reported a 60% loss probability for the SiH_2 radical on the reactor walls through laser-induced fluorescence. The experimentally obtained reaction probability falls within the range for the sticking coefficients on the H-terminated and amorphous film surfaces as determined by our atomistic calculations. Molecular-dynamics simulations of a-Si:H film growth by repeated impingement of SiH_2 radicals have revealed adsorption reactions at early stages to occur with similar energetics as the corresponding reactions of isolated radicals on crystalline surfaces. The reaction probability of

SiH2 on a-Si:H films deposited through MD simulations is approximately 30%. Finally, it is found that the SiH2 radical is much more mobile on surfaces of a-Si:H films than crystalline surfaces, especially when the hydrogen concentration in the amorphous film and, thus, on the surface is high.

SixHy clusters of various sizes have been observed in various plasmas using plasma diagnostics such as mass spectrometry. The experimental literature has assumed and employed a unit reaction probability for the clusters on all Si surfaces in phenomenological models for obtaining a-Si:H film growth rates. In addition, it is widely speculated that clusters have a detrimental effect on the film quality. To the best of our knowledge, it is still uncertain what the actual reaction probabilities are of SixHy clusters with Si surfaces, and whether and how clusters contribute towards degradation of the deposited film. We have implemented a study to calculate the reaction probability of clusters with Si surfaces, elucidate the cluster-surface interaction mechanisms and determine how they might contribute to the origin of defects such as dangling bonds and voids, etc.

A comprehensive study was also implemented of the interactions of SiH radicals originating in silane containing plasmas with crystalline and amorphous silicon surfaces based on a detailed atomic-scale analysis. The hydrogen concentration on the surface is established to be the main factor that controls both the surface reaction mechanism and the reaction probability; other important factors include the location of impingement of the radical on the surface, as well as the molecular orientation of the radical with respect to the surface. On the ordered crystalline surfaces, the radical reacts in such a way as to maximize the number of Si-Si bonds it can form even if such bond formation requires dissociation of the radical and introduction of defects in the crystal structure. The radical is established to be fully reactive with the pristine Si(001)-(2(1) surface. This chemical reactivity is reduced significantly for the corresponding H-terminated surface with hydrogen coverage of one monolayer. SiH is found to be highly reactive with surfaces of hydrogenated amorphous silicon films, independent of radical orientation and location of impingement. Our simulations predict an average reaction probability of 95% for SiH with a-Si:H film surfaces, which is in excellent agreement with experimental data.

Within the second of the 3 activities listed above, a detailed CFD/neural network simulations of the thermal-CVD process were undertaken varying three parametric variables---reactor pressure, wafer temperature and di-silane mole fraction. Because of the strong non-linearity of the parametric dependence, a relatively large number of CFD cases were required to adequately train the net over the parametric range of interest. This involved two hundred and twenty-five cases for the purposes of training the net and additional computations for purposes of validation. Preliminary results have indicated excellent agreement of the neural-net predictions with computed validation data set. More detailed validation studies are currently underway.

II.C.iii.a.6   Publications and Presentations during FY99:
Articles that appeared in journals in the period Oct 1998 - Sep 1999:

1. S. Ramalingam, D. Maroudas, and E. S. Aydil, "Interaction of SiH Radicals from a Silane Plasma with Silicon Surfaces: An Atomic Scale Simulation Study", *Journal of Applied Physics*, vol. 84, pp. 3895-3911 (1998).

2. S. Ramalingam, D. Maroudas, E. S. Aydil, and S. P. Walch, "Abstraction of Hydrogen by SiH3 from Hydrogen-Terminated Si(001)-(2X1) Surfaces", *Surface Science*, vol. 418, L8-13 (1998).

3. S. Ramalingam, D. Maroudas, and E. S. Aydil, "Visualizing Radical-Surface Interactions in Plasma Deposition Processes: Reactivity of SiH3 Radicals with Silicon Surfaces"*, IEEE Transactions on Plasma Science*, vol. 27, no. 1, 104-105 (1999).

4. S. Ramalingam, D. Maroudas, and E. S. Aydil, "Atomistic Simulation Study of the Interactions of SiH3 Radicals with Silicon Surfaces ", *Journal of Applied Physics*, vol. 86, no. 5, 2872-2888 (1999).

5. S. Ramalingam, D. Maroudas, and E. S. Aydil, "Theoretical Study of the Interactions of SiH2 Radicals with Silicon Surfaces"*, Journal of Applied Physics*, vol. 86, no.10, 5497-5508 (1999).

6. S. Ramalingam, S. Sriraman, D. C. Marra, D. Maroudas, and E. S. Aydil, "An Atomic-Scale Analysis of Deposition of a-Si:H Films on Si Substrates from SiH3 Precursor and Comparisons with Experiments", manuscript in preparation.

7. D. Maroudas, S. Ramalingam, and E. S. Aydil, "Atomic-Scale Modeling of Plasma-Surface Interactions in the PECVD of Silicon", *Symposium on Fundamental Gas-Phase and Surface Chemistry of Vapor-Phase Synthesis, The 194th Meeting of the Electrochemical Society*, November 1998, Boston, Massachusetts (invited).

8. E. S. Aydil, D. Maroudas, S. Ramalingam, D.C. Marra, B. F. Hanyaloglu, W. M. M. Kessels, P. Mahalingam, S. P. Walch, "Combined Experimental and Theoretical Study of the Radical-Surface Interactions During Plasma-Enhanced Chemical Vapor Deposition of Si films From Silane/Hydrogen Discharges", *ACS Conference* in New Orleans, August 1999 (invited).

9. D. C. Marra, S. Ramalingam, E. Edelberg, D. Maroudas, and E. S. Aydil*,* "Surface Reactions and Hydrogen Coverage on Plasma Deposited Hydrogenated Amorphous and Nanocrystalline Silicon Surfaces", *45th Annual Symposium of the AVS*, November 1998, Baltimore, Maryland.

10. S. Ramalingam, D. Maroudas, and E. S. Aydil, "Atomistic Simulation of Plasma-Enhanced Chemical Vapor Deposition of Hydrogenated Amorphous Silicon Films", *45th Annual Symposium of the AVS*, November 1998, Baltimore, Maryland.

11. S. Ramalingam, D. Maroudas, and E. S. Aydil, "Atomic-Scale Analysis of Surface Reaction Mechanisms during Plasma Deposition from Silane Discharges on Silicon Substrates", *AIChE Annual Meeting*, November 1998, Miami Beach, Florida.

12. S. Ramalingam, E. S. Aydil, and D. Maroudas, "Interaction Mechanisms of SiH3 on Si Surfaces: An Atomic Scale Simulation Study", *APS Centennial Meeting*, March 1999, Atlanta, Georgia.

13. S. Ramalingam, D. Maroudas, and E. S. Aydil, "Plasma-Surface Interactions in the PECVD of Silicon: An Atomic Scale Analysis", *MRS Spring Meeting*, April 1999, San Francisco, California.

14. Shyam Ramalingam, Dimitrios Maroudas, Eray S. Aydil, and W. M. M. Kessels "Theoretical Analysis of Reactions of SiHx Radicals and Chemically Reactive Clusters from Silane Plasmas with Crystalline and Amorphous Silicon Surfaces", To be presented at the *AVS Conference* in Seattle, October 1999.

15. Shyam Ramalingam, Dimitrios Maroudas, Eray S. Aydil, and Stephen P. Walch, "Atomistic Simulations of Radical-Surface Interactions during Plasma-Enhanced Chemical Vapor Deposition of Si Films from Silane/Hydrogen Discharges", To be presented at the *AVS Conference* in Seattle, October 1999.

16. Shyam Ramalingam, Pushpa Mahalingam, Dimitrios Maroudas, and Eray S. Aydil, "Atomistic Simulations of Radical-Surface Interactions during Plasma-Enhanced Chemical Vapor Deposition of Si Films from Silane/Hydrogen Discharges", To be presented at the *AIChE Conference* in Dallas, November 1999.

17. Shyam Ramalingam, Dimitrios Maroudas, and Eray S. Aydil, "Theoretical Analysis of Reactions of SiHx Radicals and Chemically Reactive Clusters from Silane Plasmas with Crystalline and Amorphous Silicon Surfaces", To be presented at the *AIChE Conference* in Dallas, November 1999.

18. Venkateswaran, S., Govindan, T. R., Rai, M. and Meyyappan, M., *Electrochemical Society Meeting*, Seattle, WA, May 1999.

19. Venkateswaran, S., Rai, M., Govindan, T. R., and Meyyappan, M., *SPIE Meeting*, San Jose, CA, January 2000 (to be presented).

II.C.iii.a.7   <u>Future Plans:</u>
We need to carry on the film growth simulations longer in order to build a database of the energetics and rates of the important reactions during deposition. Once a database is compiled it will be possible to develop and implement kinetic Monte Carlo techniques to simulate film growth over longer and more realistic time scales.

The named RIACS staff, S. Ramalingam, intends to graduate in the next couple of months. Another graduate student at UCSB will continue the research. He will likely analyze more reactions identified through brute-force Molecular Dynamics and work on the development and implementation of kinetic Monte Carlo techniques.

II.C.iii.a.8   Problems Encountered:
None

II.C.iii.a.9   Problem Resolution:
None needed.

II.C.iii.b   Mechanisms and Energetics of Silyl Adsorption on Silicon Surfaces

[As part of a general task in Computational and Theoretical Materials Sciences]

II.C.iii.b.1   Task Summary:
In collaboration with Dr. S. P. Walch of NASA's Ames Research Center, we conducted a comprehensive theoretical and computational study of the adsorption mechanisms and energetics of the silyl radical (SiH_3) on the dimerized Si(001) surface. The interactions of the radical with both the pristine surface and a surface terminated with one monolayer of hydrogen atoms were examined following a hierarchical computational approach that combined molecular-dynamics simulations with ab initio density functional theory calculations. The structures and energies of the equilibrium and transition-state configurations were computed for all the possible mechanisms and the detailed energetics were determined along the corresponding surface reaction pathways.

II.C.iii.b.2   RIACS Staff:
Shyam Ramalingam
Dimitrios Maroudas

II.C.iii.b.3   Overview:
The work conducted during this visit is part of our research program in the area of computational surface science focusing on quantitative analysis of plasma-surface interactions and predictive modeling of plasma deposition of silicon thin films. This is a collaborative research effort between our group at UCSB and the Computational Chemistry group in the Integrated Product Team (IPT) on Devices and Nanotechnology managed by Dr. M. Meyyappan at NASA Ames Research Center. The IPT conducts basic research in nanotechnology, computational semiconductor device physics, computational and experimental chemistry in materials processing and process/equipment modeling. The IPT's vision is to develop novel concepts in nanotechnology for NASA's future needs on electronics, computing, sensors, and advanced system miniaturization; to develop highly integrated and intelligent simulation environment that facilitates the rapid development and validation of future-generation electronic devices, as well as associated materials and processes through virtual prototyping at multiple levels of fidelity.

II.C.iii.b.4   Project Description:
This research project is based on a novel multi-scale computational approach that combines the state of the art in computational quantum and statistical mechanics. It aims at identification and analysis of chemical reactions between chemically reactive radicals from silane containing plasmas and silicon surfaces that occur during plasma deposition of silicon thin films. The approach combines molecular-dynamics, molecular-statics, and Monte Carlo simulations based on semi-empirical descriptions of interatomic interactions with ab initio quantum mechanical calculations within the framework of density functional theory. These theoretical results will provide an

accurate chemical reaction database that can be used for efficient kinetic modeling of plasma deposition over coarser laboratory time scales through kinetic Monte Carlo simulations.

A general set of interrelated tasks is described below. These tasks are pursued using the computational approach outlined above. The description of the implementation of our computational framework has been presented in our journal publications reported for FY99.

- Reaction Identification:

Chemical reactions that occur during film growth are identified through molecular-dynamics simulations over a wide range of coverage.

- Chemical Reaction Database Development:

Detailed analysis of the energetics, reaction paths, and dynamics of the reactions identified during growth. Important quantities include thermal activation energy barriers, reaction energies, detailed reaction paths for reaction rate computations, reaction probabilities of Si surfaces with SiH_x radicals (x=1,2,3) that can be used in reactor-scale modeling, and determination of deposition precursors.

- Understanding of growth mechanisms and effects of processing conditions:

The elementary processes are analyzed that lead to film deposition, hydrogen incorporation, and defect generation. Reaction mechanisms are identified and their effects on the structural quality of the film are analyzed. The mobility of radicals on the growth surfaces is analyzed. The effects on the structure and composition of the deposited films are examined of deposition parameters, such as the substrate temperature, and of minority plasma species, such as larger clusters or energetic species.

- Computational Characterization of Deposited Films:

Films deposited on crystalline substrates by molecular-dynamics simulation are characterized structurally (e.g., defect type and density, surface roughness) and chemically (e.g., hydrogen content and distribution) and compared with experimental measurements.

II.C.iii.b.5   Accomplishments during FY99:

(This refers only to the week of 8/29/99 - 9/3/99 of Professor Maroudas' visit to RIACS and NASA Ames Research Center)

The mechanism and energetics of SiH_3 adsorption on the H-terminated Si(001)-(2x1) surface was studied quantitatively through an insertion reaction that involves breaking of the Si-Si surface dimer bond. Our theoretical calculations revealed that SiH_3 adsorbs on the surface dissociatively by breaking the Si-Si dimer bond and, subsequently, transferring one of its H atoms to the H-terminated surface. As a result, two surface dihydride (SiH_2) species are generated. In conjunction with other reaction mechanisms that can form SiH_2 species on the growth surface of hydrogenated amorphous silicon films, this insertion reaction can explain the high concentrations of surface dihydride species that have been measured in plasma-deposited films.

In addition, the mechanisms and energetics of SiH_3 adsorption on the pristine Si(001)-(2x1) surface were studied quantitatively. Our theoretical calculations revealed adsorption mechanisms that include radical attachment to a surface dangling bond, insertion between surface dimer atoms accompanied by dimer bond breaking, as well as bridging of both neighboring dimer pairs in the same dimer row and adjacent dimer rows. The first two of these mechanisms are highly exothermic and are expected to be important steps in the plasma deposition process of silicon thin films following the abstraction of surface hydrogen from an originally H-terminated surface.

These analyses contributed significantly to our understanding of the chemical reactivity of SiH_3 with Si surfaces; SiH_3 is believed to be the most important precursor for the plasma deposition of Si films. For example, the reasons for the different surface hydride concentrations has been a puzzle in the literature for about two decades. Our work is offerring the first truly quantitative explanation toward a final answer to this decade-long debate.

II.C.iii.b.6   Publications and Presentations during FY99:

(Shyam Ramalingam and Professor Maroudas have reported a complete list of publications and presentations in the Project Annual Reports. In the present list, two journal papers are reported that were prepared largely during Prof. Maroudas' visit to NASA Ames Research Center and RIACS for the week 8/29/99 - 9/3/99).

1. S. P. Walch, S. Ramalingam, E. S. Aydil, and D. Maroudas, "Mechanism and Energetics of Dissociative Adsorption of SiH_3 on the Hydrogen-Terminated Si(001)-(2x1) Surface", manuscript in preparation to be submitted to *Chemical Physics Letters* [Submission expected by end of February, 2000].

2. S. P. Walch, S. Ramalingam, E. S. Aydil, and D. Maroudas, "Mechanism and Energetics of SiH_3 Adsorption on the Pristine Si(001)-(2x1) Surface", manuscript in preparation to be submitted to *Chemical Physics Letters* [Submission expected by end of February, 2000].

II.C.iii.b.7   Future Plans:

Our plans in this general area for FY 2000 have been presented in detail in Professor Maroudas Project Annual Report and Proposal for FY 2000. Our plans include at least three weekly visits by Professor Maroudas to NASA Ames Research Center and RIACS for FY 2000. These visits will focus on collaborative computational research of semiconductor surface chemical reactions with Dr. S. P. Walch and other computational chemistry researchers in the Integrated Product Team (IPT) on Devices and Nanotechnology managed by Dr. M. Meyyappan. As in FY99, Prof. Maroudas' visits will focus on conducting state-of-the-art computational work on surface chemical reaction analysis and publishing the results in leading Chemical Physics and Surface Science journals.

II.C.iii.b.8   Problems Encountered:

None

II.C.iii.b.9   Problem Resolution:

No problems

II.C.iii.c   Mechanisms and Energetics of Hydrogen Abstraction by Hydrogen from Hydrogen-Terminated Silicon Surfaces

[As part of a general task in Computational and Theoretical Materials Sciences]

II.C.iii.c.1   Task Summary:
In collaboration with Dr. S. P. Walch of NASA's Ames Research Center, we conducted a systematic theoretical and computational study of the mechanisms and energetics of hydrogen abstraction from the hydrogen-terminated dimerized Si(001) surface by H atoms impinging from a gas phase. These abstraction reactions were analyzed following a hierarchical computational approach that combined molecular-dynamics simulations with ab initio density functional theory calculations. For H atoms impinging at thermal energies, a new two-step abstraction mechanism was identified involving surface dihydride formation and subsequent desorption of a hydrogen molecule from the surface; at higher kinetic energies of the impinging H atoms, an athermal Eley-Rideal abstraction mechanism was found to be operative.

II.C.iii.c.2   RIACS Staff:
Shyam Ramalingam
Dimitrios Maroudas

II.C.iii.c.3   Overview:
The work conducted during this visit is part of our research program in the area of computational surface science focusing on quantitative analysis of plasma-surface interactions and predictive modeling of plasma deposition of silicon thin films. This is a collaborative research effort between our group at UCSB and the Computational Chemistry group in the Integrated Product Team (IPT) on Devices and Nanotechnology managed by Dr. M. Meyyappan at NASA Ames Research Center. The IPT conducts basic research in nanotechnology, computational semiconductor device physics, computational and experimental chemistry in materials processing and process/equipment modeling. The IPT's vision is to develop novel concepts in nanotechnology for NASA's future needs on electronics, computing, sensors, and advanced system miniaturization; to develop highly integrated and intelligent simulation environment that facilitates the rapid development and validation of future-generation electronic devices, as well as associated materials and processes through virtual prototyping at multiple levels of fidelity.

II.C.iii.c.4   Project Description:
This research project is based on a novel multi-scale computational approach that combines the state of the art in computational quantum and statistical mechanics. It aims at identification and analysis of chemical reactions between chemically reactive radicals from silane containing plasmas and silicon surfaces that occur during plasma deposition of silicon thin films. The approach combines molecular-dynamics, molecular-statics, and Monte Carlo simulations based on semi-empirical descriptions of interatomic interactions with ab initio quantum mechanical calculations within the framework of density functional theory. These theoretical results will provide an

accurate chemical reaction database that can be used for efficient kinetic modeling of plasma deposition over coarser laboratory time scales through kinetic Monte Carlo simulations.

A general set of interrelated tasks is described below. These tasks are pursued using the computational approach outlined above. The description of the implementation of our computational framework has been presented in our journal publications reported for FY99.

- Reaction Identification:

Chemical reactions that occur during film growth are identified through molecular-dynamics simulations over a wide range of coverages.

- Chemical Reaction Database Development:

Detailed analysis of the energetics, reaction paths, and dynamics of the reactions identified during growth. Important quantities include thermal activation energy barriers, reaction energies, detailed reaction paths for reaction rate computations, reaction probabilities of Si surfaces with SiH_x radicals (x=1,2,3) that can be used in reactor-scale modeling, and determination of deposition precursors.

- Understanding of growth mechanisms and effects of processing conditions:

The elementary processes are analyzed that lead to film deposition, hydrogen incorporation, and defect generation. Reaction mechanisms are identified and their effects on the structural quality of the film are analyzed. The mobility of radicals on the growth surfaces is analyzed. The effects on the structure and composition of the deposited films are examined of deposition parameters, such as the substrate temperature, and of minority plasma species, such as larger clusters or energetic species.

- Computational Characterization of Deposited Films:

Films deposited on crystalline substrates by molecular-dynamics simulation are characterized structurally (e.g., defect type and density, surface roughness) and chemically (e.g., hydrogen content and distribution) and compared with experimental measurements.

II.C.iii.c.5   Accomplishments during FY99:

(This refers only to the week of 9/19/99 - 9/24/99 of Professor Maroudas' visit to RIACS and NASA Ames Research Center)

A systematic study was conducted of surface reactions by which H atoms impinging from a silane-containing gas discharge on a H-terminated Si(001)-(2x1) surface abstract H atoms from the surface and return to the gas phase as H_2 molecules, while generating Si dangling bonds on the surface. The reaction mechanisms were studied by classical molecular-dynamics simulations of high-flux H impingement on the hydrogen-terminated surface of a Si(001) substrate maintained at 773 K and energetic analysis of the molecular-dynamics trajectories. Ab initio calculations within density functional theory also were carried out for accurate quantitative analysis of the abstraction reaction energetics.

Two hydrogen abstraction mechanisms were identified depending on the kinetic energy of the impinging H atom. At thermal energies, a new two-step indirect abstraction mechanism is operative. The first step consists of the barrierless formation of a surface dihydride as an intermediate reaction product followed by the thermally activated desorption from the surface of a diatomic hydrogen molecule. At higher impinging H kinetic energies, a barrierless Eley-Rideal abstraction mechanism is identified in agreement with previous proposals based on experimental studies. The thermal activation energy barriers, reaction energies, and the corresponding detailed reaction paths were computed for all the reaction steps and mechanisms that were identified.

These analyses contributed significantly to our understanding of hydrogen abstraction from hydrogen-terminated silicon surfaces, continuing a very successful study of this class of surface reactions that we had initiated in FY98. These are very important reactions for silicon film growth because they create highly reactive sites on the deposition surface. Our theoretical work offers some comprehensive quantitative explanations to decade-long debates about the possible role and importance of hydrogen abstraction reactions in the plasma deposition process.

II.C.iii.c.6 **Publications and Presentations during FY99:**
(Shyam Ramalingam and Professor Maroudas have reported a complete list of publications and presentations in the Project Annual Reports. In the present list, two journal papers are reported that were prepared largely during Prof. Maroudas' visit to NASA Ames Research Center and RIACS for the week 9/19/99 – 9/24/99).

S. Ramalingam, S. P. Walch, P. Mahalingam, E. S. Aydil, and D. Maroudas, "Mechanisms and Energetics of Hydrogen Abstraction by Hydrogen from Hydrogen-Terminated Si(001)-(2x1) Surfaces", manuscript in preparation to be submitted to *Surface Science* [Submission expected by end of February 2000].

II.C.iii.c.7 **Future Plans:**
Our plans in this general area for FY 2000 have been presented in detail in Professor Maroudas Project Annual Report and Proposal for FY 2000. Our plans include at least three weekly visits by Professor Maroudas to NASA Ames Research Center and RIACS for FY 2000. These visits will focus on collaborative computational research of semiconductor surface chemical reactions with Dr. S. P. Walch and other computational chemistry researchers in the Integrated Product Team (IPT) on Devices and Nanotechnology managed by Dr. M. Meyyappan. As in FY99, Prof. Maroudas' visits will focus on conducting state-of-the-art computational work on surface chemical reaction analysis and publishing the results in leading Chemical Physics and Surface Science journals.

II.C.iii.c.8 **Problems Encountered:**
None

II.C.iii.c.9 **Problem Resolution:**
No problems

II.C.iii.d   Computer Modeling of Deposition of Si Films using PECVD

II.C.iii.d.1   Task Summary:

A computational analysis is implemented of growth of a-Si:H films on H-terminated Si (001)-(2(1) substrate in the temperature range 300 K (T (773 K through molecular-dynamics simulations of SiHx (x=1,2,3) radical impingement. The key reactions during film growth are identified through visualization of atomic trajectories and their energetics obtained through energy minimization techniques coupled with ab initio calculations based on density functional theory. The deposited films are structurally and chemically characterized and compared with available experimental results.

II.C.iii.d.2   RIACS Staff:

Shyam Ramalingam

II.C.iii.d.3   Overview:

The work on modeling of plasma deposition of Si thin films is conducted in collaboration with the Computational Chemistry group in the Integrated Product Team (IPT) on Devices and Nanotechnology managed by Dr. Meyya Meyyappan.

The IPT at NASA Ames has been set up in the Fall of 1996 to conduct basic research in nanotechnology computational semiconductor device physics computational and experimental Chemistry in materials processing and Process / Equipment modeling. The IPT vision is to develop novel concepts in nanotechnology for NASA's future needs on electronics, computing, sensors, and advanced miniaturization of all systems; to develop highly integrated and intelligent simulation environment that facilitates the rapid development and validation of future generation electronic devices as well as associated materials and processes through virtual prototyping at multiple levels of fidelity.

II.C.iii.d.4   Project Description:

- Reaction identification: What reactions happen during film growth at sub-monolayer and multi-layer coverages?
- Database development for reactor-scale modeling.
- Reaction probabilities of SiHx (x=1,2,3) radicals with Si surfaces.
- Determine possible deposition precursors.
- Mechanistic understanding.
- Reaction mechanisms leading to good/poor quality films.
- Surface mobility of radicals on Si surfaces.
- Elementary processes that lead to film deposition, H incorporation, and defect generation.
- Effect of minority plasma species such as clusters and energetic species on deposited films.
- Effect of substrate temperature.
- Film characterization: Structural and chemical characterization of deposited films and comparisons with experiments.

II.C.iii.d.5   Accomplishments during FY99:

A computational analysis was of growth of a-Si:H films on H-terminated Si (001)-(2(1) substrate in the temperature range 500 K (T (773 K through molecular-dynamics simulations of repeated silyl-radical impingement. The key reactions during film growth were identified through visualization of atomic trajectories and their energetics obtained through energy minimization techniques coupled with ab initio calculations based on density functional theory. Reaction mechanisms were identified leading to abstraction of hydrogen atoms from the surface through Eley-Rideal and desorption mechanisms, adsorption of radicals onto the deposition surface accompanied by breaking of surface Si-Si bonds, migration of adsorbed surface species, H transfer between surface Si atoms, and formation of clusters such as disilane through reaction of adsorbed surface species and their subsequent desorption. The energetics of the important reactions were analyzed and seen to be in agreement with ab initio calculations. The deposited films were structurally and chemically characterized and the surface hydride distribution is in good agreement with corresponding experimentally grown films. The same reactions are seen to occur over the temperature range studied, although reactions such as H abstraction and desorption of surface species occur less frequently at lower temperatures; the surface population of higher (di- and tri-) hydrides is greater at lower deposition temperatures.

Silylene (SiH_2) radicals created by electron impact dissociation of silane in reactive gas discharges can play an important role in plasma deposition of amorphous and nanocrystalline silicon thin films. We have implemented a systematic computational analysis of the interactions of SiH_2 radicals with a variety of crystalline and amorphous silicon surfaces based on atomistic simulations. The hydrogen coverage of the surface and, hence, the availability of surface dangling bonds is shown to exert the strongest influence on the radical-surface reaction mechanisms and the corresponding reaction probabilities. The SiH_2 radical reacts with unit probability on the pristine Si(001)-(21) surface which has one dangling bond per Si atom; upon reaction, the Si atom of the radical forms strong Si-Si bonds with either one or two surface Si atoms. On the H-terminated Si(001)-(21) surface, the radical is found to react with a probability of approximately 50%. The SiH_2 radical attaches to the surface either by forming two bonds with Si atoms of adjacent dimers in the same dimer row or through Si-Si bonds with one or both atoms of a surface dimer. In addition, the SiH_2 radical can attach in the trough between dimer rows, forming two Si-Si bonds with second-layer Si atoms. The energetics and dynamics of these surface reactions are analyzed in detail. A reaction probability of approximately 70% is calculated for SiH_2 radicals impinging on surfaces of hydrogenated amorphous silicon (a-Si:H) films with varying concentrations of hydrogen. Recent experimental measurements have reported a 60% loss probability for the SiH_2 radical on the reactor walls through laser-induced fluorescence. The experimentally obtained reaction probability falls within the range for the sticking coefficients on the H-terminated and amorphous film surfaces as determined by our atomistic calculations. Molecular-dynamics simulations of a-Si:H film growth by repeated impingement of SiH_2 radicals have revealed adsorption reactions at early stages to occur with similar energetics as the corresponding reactions of isolated radicals on crystalline surfaces. The reaction probability of SiH_2 on a-Si:H films deposited through MD simulations is approximately 30%. Finally, it is found that the SiH_2 radical is much more mobile on surfaces of a-Si:H films than crystalline surfaces, especially when the hydrogen concentration in the amorphous film and, thus, on the surface is high.

SixHy clusters of various sizes have been observed in various plasmas using plasma diagnostics such as mass spectrometry. The experimental literature has assumed and employed a unit reaction probability for the clusters on all Si surfaces in phenomenological models for obtaining a-Si:H film growth rates. In addition, it is widely speculated that clusters have a detrimental effect on the film quality. To the best of our knowledge, it is still uncertain what the actual reaction probabilities are of SixHy clusters with Si surfaces, and whether and how clusters contribute towards degradation of the deposited film. We have implemented a study to calculate the reaction probability of clusters with Si surfaces, elucidate the cluster-surface interaction mechanisms and determine how they might contribute to the origin of defects such as dangling bonds and voids etc.

A comprehensive study was also implemented of the interactions of SiH radicals originating in silane containing plasmas with crystalline and amorphous silicon surfaces based on a detailed atomic-scale analysis. The hydrogen concentration on the surface is established to be the main factor that controls both the surface reaction mechanism and the reaction probability; other important factors include the location of impingement of the radical on the surface, as well as the molecular orientation of the radical with respect to the surface. On the ordered crystalline surfaces, the radical reacts in such a way as to maximize the number of Si-Si bonds it can form even if such bond formation requires dissociation of the radical and introduction of defects in the crystal structure. The radical is established to be fully reactive with the pristine Si(001)-(2(1) surface. This chemical reactivity is reduced significantly for the corresponding H-terminated surface with hydrogen coverage of one monolayer. SiH is found to be highly reactive with surfaces of hydrogenated amorphous silicon films, independent of radical orientation and location of impingement. Our simulations predict an average reaction probability of 95% for SiH with a-Si:H film surfaces, which is in excellent agreement with experimental data.

II.C.iii.d.6   Publications and Presentations during FY99:
Articles that appeared in journals in the period Oct 1998 - Sep 1999:

1. S. Ramalingam, D. Maroudas, and E. S. Aydil, "Interaction of SiH Radicals from a Silane Plasma with Silicon Surfaces: An Atomic Scale Simulation Study", *Journal of Applied Physics*, vol. 84, pp. 3895-3911 (1998).

2. S. Ramalingam, D. Maroudas, E. S. Aydil, and S. P. Walch, "Abstraction of Hydrogen by SiH3 from Hydrogen-Terminated Si(001)-(2X1) Surfaces", *Surface Science*, vol. 418, L8-13 (1998).

3. S. Ramalingam, D. Maroudas, and E. S. Aydil, "Visualizing Radical-Surface Interactions in Plasma Deposition Processes: Reactivity of SiH3 Radicals with Silicon Surfaces", *IEEE Transactions on Plasma Science*, vol. 27, no. 1, 104-105 (1999).

4. S. Ramalingam, D. Maroudas, and E. S. Aydil, "Atomistic Simulation Study of the Interactions of SiH3 Radicals with Silicon Surfaces ", *Journal of Applied Physics*, vol. 86, no. 5, 2872-2888 (1999).

5. S. Ramalingam, D. Maroudas, and E. S. Aydil, "Theoretical Study of the Interactions of SiH2 Radicals with Silicon Surfaces", *Journal of Applied Physics*, vol. 86, no.10, 5497-5508 (1999).

6. S. Ramalingam, S. Sriraman, D. C. Marra, D. Maroudas, and E. S. Aydil, "An Atomic-Scale Analysis of Deposition of a-Si:H Films on Si Substrates from SiH3 Precursor and Comparisons with Experiments", manuscript in preparation.

II.C.iii.d.7   Conference Presentations:

1. D. Maroudas, S. Ramalingam, and E. S. Aydil, "Atomic-Scale Modeling of Plasma-Surface Interactions in the PECVD of Silicon", *Symposium on Fundamental Gas-Phase and Surface Chemistry of Vapor-Phase Synthesis, The 194th Meeting of the Electrochemical Society*, November 1998, Boston, Massachusetts (invited).

2. S. Aydil, D. Maroudas, S. Ramalingam, D.C. Marra, B. F. Hanyaloglu, W. M. M. Kessels, P. Mahalingam, S. P. Walch, "Combined Experimental and Theoretical Study of the Radical-Surface Interactions During Plasma-Enhanced Chemical Vapor Deposition of Si Films From Silane/Hydrogen Discharges", *ACS Conference* in New Orleans, August 1999 (invited).

3. D. C. Marra, S. Ramalingam, E. Edelberg, D. Maroudas, and E. S. Aydil, "Surface Reactions and Hydrogen Coverage on Plasma Deposited Hydrogenated Amorphous and Nanocrystalline Silicon Surfaces", *45th Annual Symposium of the AVS*, November 1998, Baltimore, Maryland.

4. S. Ramalingam, D. Maroudas, and E. S. Aydil, "Atomistic Simulation of Plasma-Enhanced Chemical Vapor Deposition of Hydrogenated Amorphous Silicon Films", *45th Annual Symposium of the AVS*, November 1998, Baltimore, Maryland.

5. S. Ramalingam, D. Maroudas, and E. S. Aydil, "Atomic-Scale Analysis of Surface Reaction Mechanisms during Plasma Deposition from Silane Discharges on Silicon Substrates", *AIChE Annual Meeting*, November 1998, Miami Beach, Florida.

6. S. Ramalingam, E. S. Aydil, and D. Maroudas, "Interaction Mechanisms of SiH3 on Si Surfaces: An Atomic Scale Simulation Study", *APS Centennial Meeting*, March 1999, Atlanta, Georgia.

7. S. Ramalingam, D. Maroudas, and E. S. Aydil, "Plasma-Surface Interactions in the PECVD of Silicon: An Atomic Scale Analysis", *MRS Spring Meeting*, April 1999, San Francisco, California.

8. Shyam Ramalingam, Dimitrios Maroudas, Eray S. Aydil, and W. M. M. Kessels "Theoretical Analysis of Reactions of SiHx Radicals and Chemically Reactive Clusters from Silane Plasmas with Crystalline and Amorphous Silicon Surfaces", to be presented at the *AVS Conference* in Seattle, October 1999.

9. Shyam Ramalingam, Dimitrios Maroudas, Eray S. Aydil, and Stephen P. Walch, "Atomistic Simulations of Radical-Surface Interactions During Plasma-Enhanced Chemical Vapor Deposition of Si Films from Silane/Hydrogen Discharges", to be presented at the *AVS Conference* in Seattle, October 1999.

10. Shyam Ramalingam, Pushpa Mahalingam, Dimitrios Maroudas, and Eray S. Aydil, "Atomistic Simulations of Radical-Surface Interactions During Plasma-Enhanced Chemical Vapor Deposition of Si Films from Silane/Hydrogen Discharges", to be presented at the *AIChE Conference* in Dallas, November 1999.

11. Shyam Ramalingam, Dimitrios Maroudas, and Eray S. Aydil, "Theoretical Analysis of Reactions of SiHx Radicals and Chemically Reactive Clusters from Silane Plasmas with Crystalline and Amorphous Silicon Surfaces", to be presented at the *AIChE Conference* in Dallas, November 1999.

II.C.iii.d.8   Future Plans:
We need to carry on the film growth simulations longer in order to build a database of the energetics and rates of the important reactions during deposition. Once a database is compiled it will be possible to develop and implement kinetic Monte Carlo techniques to simulate film growth over longer and more realistic time scales.

Indicate the intended research and results for the coming project year (Oct. 1999 through Sep. 2000)

The named RIACS staff, S. Ramalingam, intends to graduate in the next couple of months. Another graduate student at UCSB will continue the research. He will likely analyze more reactions identified through brute-force Molecular Dynamics and work on the development and implementation of kinetic Monte Carlo techniques.

II.C.iii.d.9   Problems Encountered:
None.

II.C.iii.d.10   Problem Resolution:
None needed.

II.C.III.E Development of a CFD/Neural Network Based Design Tool For Semiconductor Processing

II.C.iii.e.1   Task Summary:
The research study involves the development of a practical semiconductor reactor design and process control tool using computational fluid dynamics (CFD) analyses coupled with neural networks. The CFD analysis is used to populate the design space and train the neural net. The trained net is then used to predict and optimize the design parameters.

II.C.iii.e.2   RIACS Staff
S. Venkanteswaran

II.C.iii.e.3   Overview:
The NASA Ames IPT on Devices and Nanotechnology has goals to investigate semiconductor devices and processes suitable for meeting NASA requirements on ultrahigh performance computers, fast and low power devices, and high temperature wide band-gap materials. These devices may ultimately be sub-100nm feature-size and it is important that the processes and equipment meet the stringent demands posed by the fabrication of such small devices. The IPT has therefore initiated extensive modeling efforts of the physics of these devices as well as of the processing techniques.

At the present time, reactors for CVD and plasma processing are designed by trial and error procedures. Also, once the reactor is in place, optimal processing parameters are determined through expensive and time-consuming experimentation. The availability of reliable and practical computational models and design tools would greatly simplify these tasks, while also being cost-effective. Typically, the design process involves achieving a rapid and uniform deposition on the substrate. Complete CFD simulation of all possible designs will clearly be too CPU-intensive and expensive. Neural nets provide an attractive means of performing such design studies [8]. The present study seeks to assess the feasibility of applying coupled CFD/neural networks procedures to facilitate the design process.

The candidate CFD code solves the Navier-Stokes equations coupled to species transport equations. Finite rate chemical kinetics are used to describe the chemical reactions in the gas-phase. In addition, surface reactions that occur on the substrate surface are included in the specification of the boundary condition on the substrate. For the neural net, we employ a three-layer, feed-forward network to model the functional behavior of the deposition rate as a function of the controlling parametric variables.

II.C.iii.e.4   Project Description:
The primary objective of the task is to assess the feasibility of CFD/neural network-based design tool for reactor design and process control. CFD simulations are used to populate the design space and train the neural net. The neural net is then used as a function evaluator to predict, evaluate and optimize designs. These initial investigations have focused on thermal-CVD process reactors involving di-silane. Three parametric variables---namely, the reactor pressure, wafer temperature and di-silane mole fraction---were chosen for the study and varied within prescribed parametric bounds. The neural net is used to optimize the choice of operating variables for rapid and uniform deposition of silicon on the wafer substrate.

II.C.iii.e.5   Accomplishments During FY99:
Detailed CFD/neural network simulations of the thermal-CVD process have been undertaken varying three parametric variables---reactor pressure, wafer temperature and di-silane mole fraction. Because of the strong non-linearity of the parametric dependence, a relatively large number of CFD cases were required to adequately train the net over the parametric range of interest. This involved two hundred and twenty five cases for the purposes of training the net and

additional computations for purposes of validation. Preliminary results have indicated excellent agreement of the neural-net predictions with computed validation data set [1]. More detailed valiadation studies are currently underway [2].

II.C.iii.e.6   Publications and Presentation During FY99:

1. Venkateswaran, S., Govindan, T. R., Rai, M. and Meyyappan, M., *Electrochemical Society Meeting*, Seattle, WA, May 1999.

2. Venkateswaran, S., Rai, M., Govindan, T. R., and Meyyappan, M., *SPIE Meeting*, San Jose, CA, January 2000 (to be presented).

II.C.iii.e.7   Problems Encountered:

The large number of CFD cases necessitated by the strong non-linearity in the parametric dependence was not initially anticipated. Further, the accurate prediction of the silicon deposition rate required much tighter convergence tolerances than are typically used in CFD. This is because the deposition rate is determined by gradients at the wafer surface, which converge more slowly than the rest of the flowfield. Both the above factors resulted in greater time requirements for generating the CFD dataset.

II.C.iii.e.8   Problem Resolution:

The computation of the large number of CFD cases was accomplished by running them in parallel (typically in sets of 25) on the Origin 2000 cluster at NASA Ames.

II.C.iv   Computational Fluid Dynamics

Several visiting scientists worked at RIACS in the CFD area. The following are their reports.

II.C.iv.a   Advanced Computation Methods for Compressible Low Mach Number Flows

II.C.iv.a.1   Task Summary:

The development of high-order finite difference methods for computational aeroacoustics is progressing as planned.

II.C.iv.a.2   RIACS Staff:

Dr. Bernhard Müller, RIACS visitor from 12 to 23 July 1999,
B. Müller's visit at RIACS was financed by a grant of the Swedish Research Council for Engineering Sciences.

II.C.iv.a.3   Overview:

Computational aeroacoustics has been emerging as a new discipline due to the increasing demand for noise reduction in aerodynamics. The design principle for low dissipative efficient high order numerical methods for computational aeroacoustics is very different than for standard transonic or low speed computational fluid dynamics since high accuracy has to be maintained over large distances and times. The development, analysis and application of high-order schemes in aeroacoustics are challenging.

II.C.iv.a.4   Project Description:
Within the project "Accurate Numerical Simulation of Vortex Sound at Low Mach Numbers", a high-order finite difference method will be developed, verified and applied to improve the physical understanding of vortex sound and its control.

II.C.iv.a.5   Accomplishments during FY99:
During the two weeks stay at RIACS, B. Müller collaborated with H.C. Yee and extended his 2nd-order 2D Euler code to include the recently developed Yee et al. high order low dissipative schemes employing artificial compression method filters. The Yee et al. schemes are suitable for long time integration. The effort also included Müller's perturbation formulation suitable for low Mach numbers. The new method was verified for a radial pressure pulse.

II.C.iv.a.6   Publications and Presentations during FY99:
None. Due to the shortness of the visit, there was not enough time to reach a stage that resulted in a publication. More work is required. Perhaps a subsequent visit will help the research reach up to a stopping point

II.C.iv.a.7   Problems Encountered:
Preliminary results for the numerical simulation of sound emitted by an elliptic vortex patch, where B. Müller provided the analytical solution, indicate numerical problems with the boundary conditions.

II.C.iv.a.8   Problem Resolution:
We plan to employ boundary conditions satisfying the summation by parts property developed at the Department of Scientific Computing of Uppsala University. The use of entropy splitting in conjunction with the Yee et al. schemes can help minimize the use of numerical dissipation.

II.C.iv.b   Computational Fluid Dynamics Parallel

II.C.iv.b.1   Task Summary:
A new parallel scalable flow solver for in viscid flow is being developed for adaptively refined Cartesian meshes with embedded boundaries. It includes an on-the-fly parallel domain decomposition based on space filling curves.

II.C.iv.b.2   RIACS Staff:
Marsha Berger, RIACS Visiting Fellow

II.C.iv.b.3   Overview:
The use of Cartesian meshes with embedded boundaries has greatly increased in the last 5 to 10 years. This is due to their ability to handle extremely complicated geometry in an automatic and robust way. Mesh generation is now so fast and efficient that flow solution time is again becoming a bottleneck. Thus, we are developing a new flow solver, where we pay particular attention to parallel scalability and multi-grid convergence acceleration.

II.C.iv.b.4   Project Description:
Research objectives include:
1.   Develop highly scalable parallel partitioning algorithm.
2.   Develop numerical discretization suitable for cutcells with widely varying non-smooth cell volumes.
3.   Develop multigrid coarsening strategy suitable for non-body-fitted Cartesian grids.

II.C.iv.b.5   Accomplishments during FY99:
We have designed and implemented our parallel decomposition algorithm. Our implementation can decompose the mesh into any number of partitions at run time. It achieves speedups in excess of 52 using 64 processors. The multigrid coarsening algorithm has been developed and implemented, and achieves coarsening ratios in excess of 7 on reasonable sized three-dimensional meshes.

II.C.iv.b.6   Publications and Presentations during FY99:
"A Domain-Decomposed Multilevel Method for Adaptively Refined Cartesian Grids with Embedded Boundaries" with M.J. Aftosmis and G. Adomavicius. Accepted, *AIAA Paper 2000-0808*, Reno, NV, Jan. 2000.

"Automatic Hybrid-Cartesian Grid Generation for High-Reynolds Number Flows around Complex Geometries'', with M. Delanaye, M. Aftosmis, Y. Liu and T. Pulliam, *AIAA Paper 99-0777*, Reno, NV, Jan. 1999.

II.C.iv.b.7   Future Plans:
Still to be developed are improved spatial discretizations and limiters for use on the cutcells of the Cartesian mesh that intersect the geometry. Another important open question is the level of geometry fidelity that needs to be preserved when generating coarse meshes in the multigrid hierarchy.

II.C.iv.c   Turbulence Modeling

II.C.iv.c.1   Task Summary:
This project aims to translate developments in numerical algorithms into practical tools to numerical simulation of shock wave interactions with turbulent flow. In particular the goal is to carry out such simulations with low dissipative high-order numerical schemes.

II.C.iv.c.2   RIACS Staff:
Neil Sandham (Visiting Scientist)

II.C.iv.c.3   Project Description:
This project provides support for short visits by Prof. Sandham to NASA-Ames. The aim of the work is to demonstrate practical applications of recent advances in numerical techniques relating particularly to the efficient simulation of shock wave interactions with boundary layer turbulence.

II.C.iv.c.4   Accomplishments during FY99:
Prof. Sandham visited NASA-Ames for a two-week visit from 11th to 25th August 1999. This two-week visit was devoted to the further development of a high-order compressible turbulence solver. Previous work has examined shock-capturing methods with only localized dissipation to avoid excessive damping of turbulence. The emphasis on this visit was on testing an entropy splitting approach for three-dimensional turbulence. This follows an observation in Yee et al (1999) that there may be an advantage to this approach for shock-free problems, whereas for problems with shocks the numerical dissipation dominates and there is little difference between split and unsplit approaches. A turbulent channel flow test case was devised in order to test the approach. The centerline Mach number is 2 and the Reynolds number 200. Choosing a small box allows a reduced number of modes, while retaining the main characteristics of near-wall turbulence. Further code optimizations were carried out to improve the basic performance of the code and allow later parametric studies of sensitivity to numerical parameters in realistic program run times. An early finding of the work is that with entropy splitting it is possible to run the calculation with no artificial dissipation or filtering at all, even when the resolution is marginal. Previous compressible turbulent flow calculations have had to use very fine resolutions to avoid numerical problems. Further work is planned to develop an even more robust formulation of the basic equations, with comparable performance to those known for incompressible flow simulations. Further simulations will be run in Southampton between now and the next visit, planned for January 2000.

II.C.iv.c.5   Publications and Presentations during FY99:
1. Yee HC, Sandham ND, Djomehri MJ 1999 "Low-dissipative high-order shock-capturing methods using characteristic-based filters," *Journal of Computational Physics,* Vol.150, No.1, pp.199-238

II.C.iv.c.6   Future Plans:
Further work is planned to develop an even more robust formulation of the basic equations, with comparable performance to those known for incompressible flow simulations. Further simulations will be run in Southampton between now and the next visit, planned for January 2000.

II.C.iv.c.7   Problems Encountered:
None

II.C.iv.d   Convergence Accelaration of Numerical Methods for the Navier-Stokes Equation

II.C.iv.d.1   RIACS Staff:
David Zingg
Thomas Pulliam (NASA Ames)

II.C.iv.d.2   Task Summary:
The task consisted of further study of a new higher-order spatial discretization for computing aerodynamic flows developed at the University of Toronto Institute for Aerospace Studies. Detailed comparisons with a number of modern schemes, including a third-order upwind-biased

flux-difference split scheme, demonstrate the superiority of the higher-order approach with respect to efficiency. In addition, Drs. Pulliam and Zingg completed a textbook entitled Fundamentals of Computational Fluid Dynamics, co-authored with Harvard Lomax, which has now been submitted to Springer for publication.

II.C.iv.d.3   Overview:
Drs. Pulliam and Zingg have worked together for many years developing and investigating algorithms for the compressible Navier-Stokes equations applicable to aerodynamic flows. Much of this work has centered on the NASA Ames flow solver ARC2D, which is the basis for the widely used NASA flow solver OVERFLOW. The objective is to develop algorithms, which improve the efficiency and reliability of ARC2D and OVERFLOW.

II.C.iv.d.4   Project Description:
The current project involves the development of a spatial discretization which raises all approximations to an order of accuracy consistent with third-order global accuracy. Grid convergence studies demonstrate that this algorithm produces substantial reductions in numerical error relative to the discretizations currently in ARC2D and OVERFLOW, including modern schemes such as a third-order upwind-biased flux-difference split scheme with limiting. Work is underway to implement the new algorithm in OVERFLOW.

II.C.iv.d.5   Accomplishments During Summer 99:
During this time period, grid convergence studies were carried out in order to determine the numerical errors resulting from various spatial discretizations. Several flow cases were examined, including both subsonic and transonic flows over airfoils. These studies provide a convincing demonstration of the benefits of the higher-order algorithm in the context of practical steady turbulent flows over airfoils.

II.C.iv.d.6   Related presentations:
"Comparison of Several Spatial Discretizations for the Navier-Stokes Equations," presented at the *14th AIAA Computational Fluid Dynamics Conf.*, Norfolk, June 1999.

"Recent Progress in Computational Aerodynamics at UTIAS," presented at the NASA Ames Research Center and at UC Davis, Aug. 1999.

II.C.iv.d.7   Related publications:
1. Zingg, D.W., De Rango, S., Nemec, M., and Pulliam, T.H., "Comparison of Several Spatial Discretizations for the Navier-Stokes Equations," *AIAA 99-3269*, June 1999, also submitted to *J. Comp. Phys*, Nov. 1999.

2. Zingg, D.W., and Pulliam, T.H., "A Textbook for a First Course in Computational Fluid Dynamics," *Computational Fluid Dynamics Journal*, in press (invited).

3. Lomax, H., Pulliam, T.H., and Zingg, D.W., "Fundamentals of Computational Fluid Dynamics," textbook submitted to *Springer*, Oct. 1, 1999.

II.C.iv.e   Dynamic Load Balancing for Grid Partitioning on SP-2

II.C.iv.e.1   Task Summary

I visited RIACS for 7 months (6/98-1/99) under RIACS travel support. While at RIACS, I worked on scalable parallel sorting and graph partitioning for large-scale distributed-memory multiprocessors. The main objective of the work was to help investigate the performance of distributed-memory machines.

Two sorting algorithms were used for the study including bitonic sorting and radix sorting. Two distributed-memory machines installed at the Maui High Performance Computing Center and at NERSC, Lawrence Berkeley Laboratory were used to obtain experimental results. The algorithms were parallelized to exploit the nature of the distributed-memory machines. All were done from scratch.

The sorting algorithm, called Balanced Radix Sort, is the fastest to date on large-scale distributed-memory machines. Experimental results indicated that balanced radix sort can sort 0.5G integers in 20 seconds and 128M doubles in 15 seconds on a 64-processor IBM SP2-WN while yielding over 40-fold speedup. When compared with the fastest sorting algorithms, balanced radix sort was 30% to 100% faster based on the same machine and key initialization.

The graph partitioner called, S-HARP, is the fastest dynamic graph partitioner for adaptive computations known to date. Experimental results indicate that S-HARP can partition a mesh of over 100,000 vertices into 256 partitions in 0.18 seconds on a 64-processor Cray T3E. S-HARP is three to 15 times faster than the other dynamic partitioners on computational meshes of size over 100,000 vertices while giving comparable edge cuts.

The study is ongoing and some more results will be published in the future.

II.C.iv.e.2   RIACS Staff:

Andrew Sohn

II.C.iv.e.3   Publications and Presentations during FY99:

1. Sohn and Y. Kodama, "Load Balanced Parallel Radix Sort", in *Proc. of the 12th ACM International Conference on Supercomputing*, Melbourne, Australia, July 1998, pp.305-312.

2. Sohn and H. D. Simon, S-HARP: "A Scalable Parallel Dynamic Partitioner for Adaptive Computations", in *Proc. of Supercomputing 98*, November 1998, Orlando, Florida.

II.D   Applications of Information Technology

II.D.i   Advanced visualization and collaborative virtual environments for medical and scientific imaging

II.D.i.a   Task Summary:

Development of advanced 3D reconstruction and visualization techniques for cell biology and medical research, and collaborative virtual environments for interactive data analysis. This is a multi-project task oriented towards creating virtual reality technologies associated with biological models. Projects in 1999 included: High bandwidth distributed imaging, multiresolution data display, Wide Area Telemedicine, Distributed Virtual Environments, and 3D Image segmentation (CT, MRI, and Ultrasound) and reconstruction.

II.D.i.b   RIACS Staff:

Esfandiar Bandari
Rei Cheng
Xander Twombly

II.D.i.c   Overview:

Computer science research at the Ames Center for Bioinformatics is focused primarily on the development of advanced 3D reconstruction and visualization techniques for cell biology and medical research. RIACS staff comprises a majority of the computer scientists at the Center for Bioinformatics, and provides the technical leadership for all CS projects in the center. The main project was the development of the Virtual Collaborative Clinic (VCC), a system designed for sharing complex 3D reconstructions of anatomical objects for use in wide area telemedicine. The VCC is an ongoing project to develop an interactive, collaborative virtual environment for use by physicians for diagnosis and treatment planning. The virtual environment is designed for the display and real time manipulation of 3D models generated from patient-specific volumetric data sets such as CT, MRI, and Ultrasound. Such virtual environments are used to examine surface reconstruction of organs such as the heart, lung, bone, and skin, and aid in the diagnosis of physical abnormalities and surgical treatment planning. Expansion of the basic virtual environment to a collaborative environment entails linking the activity of each individual user at their own client site to all of the other participants in the simulation, effectively creating a shared world in which the models and participants reside.

II.D.i.d   Project Description:

- Multiresolution representation and display of complex 3D surface models. Multiresolution display allows the software to automatically adjust the modeling fidelity to maintain interactive display rates in a virtual environment.
- Real-Time linkage of 3D scenegraphs in multiple virtual environments. Provides the basic connectivity between virtual environments in a collaborative setting.
- High bandwidth image replication protocol. Used to leverage the rendering capabilities of a graphics supercomputer such as an SGI Onyx2 Infinite Reality2 system.
- 2D Ultrasound segmentation for the analysis and measurement of heart valve defects.

II.D.i.e   Accomplishments during FY99:

Virtual Collaborative Clinic v1.0 was completed and field-tested in FY99. The following components were functional in the demonstration:

- Multiresolution display – two resolutions per object. The display system toggled between the low-resolution representation during manipulation of the object, and the high-resolution image for static display.
- Collaborative environment – using the World2World toolkit from Sense8 Inc. created a method for automatically analyzing and linking the objects in a scenegraph between all virtual environments participating in the collaboration.
- Remote display server – created an encoding protocol for capturing the display of the virtual environment on a graphics supercomputer and transmitting the images to all systems in the collaboration via the IP multicast protocol. This system stressed the high bandwidth capabilities of the next generation Internet and contained tuning control to account for network bottlenecks.
- Developed a segmentation method for use with Ultrasound data based on a modification of the Metropolis algorithm to aid in the identification of heart valve defects. Subsequent images were used in the Virtual Collaborative Clinic demonstration.

II.D.i.f Publications and Presentations during FY99:

- Public demonstration of the VCC system on May 4[th], 1999. The program was used to link physicians at ARC, Stanford University Hospital, Salinas Valley Memorial Hospital, Cleveland Heart Clinic, and the Navajo Nation health clinic at Shiprock, New Mexico via satellite.
- Awarded NASA Software of the Year runner-up, 1999.
- Submission of 3 patient applications for software components used in the VCC system.

II.D.i.g Future Plans:

The focus of the Center for Bioinformatics is shifting from straight visualization to a more interactive technology. Leveraging the previously developed virtual reality environments, we are beginning to develop a simulation engine for real-time interactions with dynamic 3D models. In particular, we are developing a finite element-based simulator of soft tissue dynamics for use in simulated surgery. The surgical tools will be used in the development of a virtual glovebox (the glovebox workstations to be used on the International Space Station for biological experiments), designed to aid in the training of astronauts in animal surgery. The virtual glovebox will comprise an immersive environment of the same physical dimensions as the real glovebox, a series of force-feedback devices to provide realistic "feel" of the objects in the box, and a variable gravity to simulate all conditions under which the box may be used (terrestrial, zero gee, and centrifuge simulated gravity).

II.D.i.h Problems Encountered:

No problems.

II.D.i.i Problem Resolution:

No Problems.

II.D.ii   Multi-institutional Collaboration on R&D, and Demos for Natural Hazards Impact Reduction

II.D.ii.a   Task Summary

This project facilitates collaborative research and development between NASA and non-NASA organizations on natural disasters and related technologies, demonstration projects, and model development and validation. This activity complements and supports research and technology development by the Earth Sciences Division (ESD) at Ames Research Center. The potential applications of remotely sensed data from NASA satellites and aircraft to mitigation of natural disasters are investigated.

II.D.ii.b   RIACS Staff

Richard G. Johnson (consultant)

II.D.ii.c   Overviews

The capabilities to acquire, process, and distribute data from multiple sources, including satellites and both piloted and un-piloted aircraft, to mitigate the scope and severity of natural disasters have continued to expand rapidly. The Earth Sciences Division at NASA Ames Research Center has conducted extensive disaster-related research and technology development, particularly on fires and floods for more than a decade. There has been rapid growth in the capabilities of information and communications systems to process large volumes of data, to extract actionable information, and to distribute and visually display the information to the user communities. These growing systems capabilities provide new opportunities to develop and implement the needed technologies to conduct the research and related technology demonstration projects to reduce the large fiscal and human impacts of natural disasters. The national costs of natural disasters from 1992-1997 were $54 billion dollars per year, with about 10% of these costs occurring in California.

II.D.ii.d   Project Description

This project reviews and recommends approaches for collaborative disaster-related programs and information exchanges between the Ames Earth Sciences Division and federal and state agencies, local and regional government, industries, not-for-profit organizations, and universities.

This project also reviews and assesses the applicability of advanced information systems and wireless communication systems for near-term and future ESD research and development activities related to natural disasters.

II.D.ii.e   Accomplishments

This activity was initiated in July 1999, and has focused primarily on identifying and structuring potential disaster-related programs for collaboration between the Ames Earth Science Division (ESD) and other non-NASA partners, including RIACS.

Briefing materials were prepared and meetings held with the ESD Chief, Deputy Chief, Ecosystems Science and Technology Branch Chief, and the RIACS Director to explore Disaster

Infosphere concepts and architectures that were focused explicitly on direct access to disaster-related information by all users. This information-centric approach is intended to provide rapid access with minimum user constraints to dynamic information that disaster manager and other users require making effective decisions before the disaster-related information degrades in value or becomes totally obsolete.

The San Jose State University (SJSU) Foundation has recently established and funded a Collaborative for Disaster Mitigation (CDM). Meetings to discuss the options for collaboration between the CDM and the Ames Research Center were held with the Branch Chief for the ESD Ecology Branch, Deputy Director of the Ames Commercial Technology Office, the SJSU President and the CDM Director. It was concluded that a collaboration program between CDM and ARC would be of significant value to both participants and should be formally established. Steps are being taken to implement a formal agreement.

The potential high-value roles for Unpiloted Air Vehicles (UAV) have been investigated for acquiring disaster-related information. Two particularly high-value roles for such platforms are their use in acquiring timely and repetitive remotely-sensed data on wildfires and their potential to serve as high bandwidth wireless communication nodes during disaster events. These applications and the needs for low-power and lightweight on-board data systems were reviewed with the Chief of the ESD Ecosystems Science and Technology Branch.

II.D.ii.f   Publications and Presentations
None

II.D.ii.g   Future Plans
The plan is to continue the project into FY00 along the same lines but at an increased scope of effort.

II.D.ii.h   Problems Encountered
None

II.D.ii.i   Problem Resolution
None needed

III   <u>Seminars</u>

Manny Rayner (RIACS), "An Overview of the Spoken Language Translator", December 13, 1999.

In this talk, I will give an overview of the Spoken Language Translator (SLT) project. SLT, which ran from 1992 to 1999 under sponsorship from Telia Research, Stockholm, was one of the first serious projects in the area of automatic translation of speech. The final prototype performed translation between English, French and Swedish in the domain of air travel inquiry systems, using a vocabulary of about 1500 words, and with an accuracy of about 75%. Average processing times are around 4 seconds per utterance. The SLT system consists mainly of general purpose domain-independent components which are tuned to the domain using supervised training techniques, and uses a hybrid architecture which combines deep and shallow processing methods. Many of the techniques developed under the project are applicable to other types of spoken language understanding task. A book about the SLT System will be published next year by Cambridge University Press.

Andreas Stolcke, SRI International, "Combining Prosodic and Language Models for Speech Segmentation and Recognition", December 02, 1999.

This talk will give an overview of recent work at SRI aimed at extending the output of speech recognizers beyond the usual stream of words, providing additional information such as sentence segmentation, disfluency detection, identification of proper names and topic segmentation. All of these tasks are addressed in a common approach: a combination of the (fairly standard) lexical modeling, using hidden Markov and statistical language modeling, and the (not-so-standard) exploitation of prosodic cues to the "hidden" events (tags) sought. I will discuss the modeling approaches used, and present results for sentence segmentation and disfluency annotation on the Switchboard corpus, as well as for sentence and topic segmentation in the Broadcast News domain. Prosodic modeling is shown to improve performance on these tasks, in some cases dramatically. It turns out that hidden event modeling can also improve word recognition itself, by constraining hypotheses to be consistent with prosodic characteristics of hidden events.

This is joint work with Liz Shriberg, Dilek Hakkani-Tur and Gokhan Tur.
Bio: Andreas Stolcke is a Sr. Research Engineer in SRI's Speech Technology and Research Laboratory, where he has been working on statistical language modeling, spontaneous speech modeling, and large-vocabulary recognition since 1994. Prior to that, he received his undergraduate education at the Technical University of Munich, and a Ph.D. from UC Berkeley, with a thesis on Bayesian unsupervised learning of grammatical structure.

Butler Hine (RIACS), "The Intelligent Systems Program", November 18, 1999.

Dr. Butler Hine has just returned to NASA Ames as a new Level 2 Program Manager for the Intelligent Systems (IS) Program. His area of responsibility will be the Autonomous Reasoning

(AR) area, the goal of which is to create smart spacecraft, rovers, and ground systems capable of achieving mission goals independent of low-level human control. Prior to returning to NASA Ames, Butler was CEO of a Silicon Valley software start-up company and before that was the head of the Intelligent Mechanisms Group in Code IC. Butler's talk will begin with a brief summary of his background followed by a discussion of the goals and organization of the Autonomous Reasoning Program within IS.

Charles Pecheur, "Model Checking for Autonomy Software", November 04, 1999.

To fulfill its mission of deep space exploration in a "faster, better, and cheaper" way, NASA is putting a lot of effort in autonomous software such as Ames' Remote Agent (RA), demonstrated last May on Deep Space One (DS-1). The validation of such complex software systems poses a big challenge: because of internal decision-taking and concurrency, the range of possible situations becomes so large that traditional black-box testing is very inefficient. Achieving and assessing the reliability of autonomous systems will require the use of advanced V&V techniques such as model checking.

After a succinct introduction to the ins and outs of model checking, this presentation will survey past, ongoing and planned work on applying model checking to autonomy software in NASA Ames' Automated Software Engineering Group. As an initial case study, our team found important concurrency bugs in the executive subsystem of the RA, using the Spin model checker. Building on initial work at CMU, we are now extending a compiler that feeds models used by the Livingstone fault recovery system into the SMV symbolic model checker. Some similar work has been performed on a model for the HSTS planner (HSTS and Livingstone are also part of the RA). The next stage will be to address the validation of a model-based system as a whole, which poses a difficult problem: such a system is built around an inference engine that would be very hard to model in a form amenable to a model checker.

Manny Rayner, "Turning Speech into Scripts", October 21,1999.

We describe an architecture for implementing spoken natural language dialogue interfaces to semi-autonomous systems, in which the central idea is to transform the input speech signal through successive levels of representation corresponding roughly to linguistic knowledge, dialogue knowledge, and domain knowledge. The final representation is an executable program in a simple scripting language equivalent to a subset of CSHELL.

At each stage of the translation process, an input is transformed into an output, producing as a byproduct a "meta-output" which describes the nature of the transformation performed. We show how consistent use of the output/meta-output distinction permits a simple and perspicuous treatment of apparently diverse topics including resolution of pronouns, correction of user misconceptions, and optimization of scripts. The methods described have been concretely realized in a prototype speech interface to a simulation of the Personal Satellite Assistant.

Hector Garcia-Molina (Stanford University), "Barriers to Digital Libraries", October 07, 1999.

In this talk, the speaker will discuss the challenges and barriers faced in implementing a modern digital library, including technical, legal, and societal ones. In particular, he will also give an overview of the Stanford Digital Library Project, and describe how we are addressing some of the technical barriers in the areas of interoperability, preservation, searching, mobile access, and protection of intellectual property.

Andreas Frank, Ph.D. , "A Transient Coupled Fluid/Solid Numerical Simulation Of The Aortic Valve For Normal And Calcified Leaflets", September 23, 1999

Aortic stenosis (AS) is an aortic valve disease causing flow obstruction during left ventricular (LV) systolic ejection due to a reduced aortic valve area (AVA). AS can be caused by calcification which limits leaflet excursion resulting in decreased AVA, increased flow resistance, and elevated LV pressures. Clinical assessment utilizes calculated AVAs, however, AVA may be a function of flow and thus not adequately discriminate between moderate and severe AS for varying flow conditions.

Numerical simulations for increased leaflet calcification (10x, 50x, and 200x normal) resulted in increased AS, as well as: 1) reduced leaflet displacement and AVA, 2) diminished slow closure phase, 3) increased transvalvular flow velocities with decreased flow, 4) increased transvalvular pressure gradients and resistance, 5) increased sinus vortex size, growth and strength, and 6) increased flow regurgitation. For severe AS with fixed calcification the resistance was nearly independent of flow, whereas the AVA was a linear function of flow (50% flow decrease provided 23% AVA decrease). Hence the resistance, not the AVA, better distinguished between moderate and severe AS for a variety of flow conditions, providing a better assessment of AS severity.

John O'Neill, Ph.D., "Mission Control: Technology Paradise or Technology Wasteland?", September 09, 1999.

The Human-Centered Computing (HCC) groups have been participating in the design of a new logging/handover tool to support the Reduced Increment Manning concept for the International Space Station's Mission Control. Reduced Increment Manning is exploring whether Mission Control can be operated without a full flight control team overnight and at weekends.

A key issue for ISS Mission Control is whether the artificial distinctions between "mission-critical" software and "PC-software" is meaningful. It can be argued that much of the information that has traditionally existed on the non-critical PC systems will actually be mission-critical in long duration missions.

The system integration challenges are more extensive than simply introducing new technologies. There are basic infrastructure issues that need to be addressed ranging from simple issues like networking the systems together, to more complex issues like how will ISS Mission Control at JSC collaborate with Mission Control at Moscow, with Mission Control at Marshall, and with Mission Control Centers for the Europeans, Japanese, and Canadians - and what technologies will enable these collaborations.

Finally, I will discuss some of the technological opportunities at Mission Control, and some of the paradigm-shifts that are preventing the introduction of new technologies.

Dr. Anil Kokaram (RIACS), "Techniques for the Automated Restoration of Degraded Film and Video: Motion estimation in the presence of missing data", August 25, 1999.

Anil Kokaram completed his Ph.D. at the Signal Processing Group of the Engineering Department, Cambridge University in 1993. He remained there as a research fellow until 1998 when he took up his present post of Lecturer in the Department of Electronic and Electrical Engineering at the University of Dublin, Trinity College, Dublin, Ireland. His main interest has been in the broad area of video processing and in problems in the restoration of degraded film and video including motion estimation for degraded sequences (e.g., Video acquired for particle image velocimetry). He has a particular interest in Bayesian techniques and MCMC for signal processing.

The increasing importance of digital video in the broadcast television and film industries has placed heavy demand on holders of television and film archives. Unfortunately, archived material is typically in very bad condition. Scratches, missing data, line jitter and shake are all typical examples of degradation possible with time. It is expected that the consumer will come to demand a better quality image as part of the "digital package". Furthermore, with the adoption of several compression schemes for broadcast, it is not sensible to waste bandwidth transmitting degraded portions of image material. Therefore, removing these artifacts has been of interest for some time. Most of the post-production houses perform much of this re-touching by hand. This is a very painful process.

This talk will introduce some of the problems in degraded video and film and present automated techniques for the treatment of the problems. Much of the material concerns the removal of missing data and an overview of a Bayesian approach to the problem will be given. In particular, this solution has important implications for the estimation of motion in video sequences in the presence of missing data. An application for robust video communications in the presence of high error rates (e.g., wireless video, internet video) will also be illustrated.

Equipment allowing, the talk will conclude with some video demos of restorations on real data.

Dr. Geoffrey Briggs, Code S, "CMEX experiments with the University of West Florida Institute for Human and Machine Cognition's (UWFa-IHMC) Concept Map Navigator for

educational outreach and for collaborative research: a precursor to more general Astrobiology applications", May 21, 1999.

Dr. Geoffrey Briggs, Scientific Director of NASA Ames' Center for Mars Exploration (CMEX), will speak about an application of concept maps for navigating voluminous Mars data collections. Prior to his current position, Dr. Briggs was Director of NASA's Solar System Exploration Division in the Office of Space Science and Applications at NASA HQ. He was at HQ during the launch of the Pioneer Venus spacecraft, the Pioneer 11 encounter with Saturn, the launch of Galileo, Ulysses and Magellan and the Voyager launches, and the start of Magellan, Mars Observer and Cassini missions. Before serving at NASA HQ, Dr. Briggs worked at JPL and was a science investigator on the Mariner 9, Viking Orbiter and Voyager Imaging Teams.

The Center for Mars Exploration (CMEX) is developing a new version of its teacher resource CD-ROM with a browser (developed by UWFa-IHMC) that uses concept maps to aid the user to better understand, and to more easily navigate within, the voluminous content of the CD. I will report on progress in the ongoing development of the CD in collaboration with the IHMC.

CMEX is also preparing to use the Concept Map browser as part of its ongoing support for the Mars Surveyor Project Office in selecting landing sites for the upcoming lander/rover/sample return missions. Ames' CMEX is involved because astrobiology objectives are central to the missions' rationale. We (Virginia Gulick and others) are organizing annual science workshops to review current Mars mapping and analysis by the science community and, in support of these workshops, have established a web site that we intend to be state-of-the art for collaborative research. To this end CMEX is in collaboration with Code IN (Glenn Deardorff and others) and we will report progress in developing the collaborative web site. It is hoped that this site will be a prototype for collaborative research within the Astrobiology Institute and the astrobiology community in general. Jeff Scargle will comment on the potential."

Dr. P. Pandurang Nayak, RIACS; Dr. Peter Norvig (Code IC),"The New Millennium Remote Agent: To Boldly Go Where No AI System Has Gone Before", May 07, 1999

The New Millennium Remote Agent is an autonomous spacecraft control system being developed jointly by NASA Ames and JPL. It integrates constraint-based planning and scheduling, robust multi-threaded execution, model-based diagnosis and reconfiguration, and real-time monitoring and control. The Remote Agent will control Deep Space One (DS-1), the first of NASA's New Millennium missions launching in late 1998. As the first AI system to autonomously control an actual spacecraft, the Remote Agent will enable the establishment of a "virtual presence" in space through an armada of intelligent space probes that autonomously explore the nooks and crannies of the solar system. In this talk I will describe the Remote Agent architecture and its main components, with a special focus on the model-based diagnosis and reconfiguration system embodied in the Livingstone program. Livingstone is a kernel of a reactive model-based autonomous system that performs significant deduction within the reactive control loop.

Based on the IJCAI-97 Invited Talk given jointly with Nicola Muscettola, Barney Pell, and Brian C. Williams.

Dr. Yuri Gawdiak (Code IC); Dr. Beth Ann Hockey (RIACS); Dr. Frankie James (RIACS), "RIACS Language Interfaces and Speech Technology Group (RIALIST)", April 30, 1999.

Beth Ann Hockey and Dr. Frankie James, members of the RIACS core research project, will present their plans and talk about the current status of their work in cooperation with the Personal Satellite Assistant (PSA) and Surface Movement Advisor (SMA) projects. Later in the afternoon Dr. Hockey and Dr. James will host a demonstration at their RIACS office. The cooperative effort will be introduced by NASA scientist Yuri Gawdiak, Level 2 Manager for the Aviation Safety Program, lead PI for the PSA project and project manager for SMA.

There are many situations in which a standard desktop computer interface is difficult, dangerous, or impossible to use. For example, if the user is performing a task that requires their complete visual attention, a GUI display can be distracting. Similarly, if the user is wearing gloves or needs to use their hands to perform an experiment, keyboard input is tricky, at best. In these cases, other interface technologies, especially speech, can alleviate the problems faced by the users.

Speech interfaces are becoming increasingly more common in desktop computing, with the advent of affordable dictation software by companies such as Dragon Systems and IBM. Flexible command and control applications, which are of the greatest potential use to NASA, are not the focus of commercial development. Since these systems are not designed with NASA in mind, they are therefore (1) not designed for use in physically challenging environments, and (2) lack specific technical language models for space and aeronautics domains.

In this talk, we will discuss areas of speech technology of particular interest to NASA, including areas of basic research that will extend the capabilities of speech applications for supporting future NASA projects. As part of this discussion, we will outline our speech interface development for two current NASA projects: the Surface Movement Advisor (SMA) and the Personal Satellite Assistant (PSA).

Michael Shafto (Human-Automation Integration Research Branch); Dr. William Clancey (Chief Scientist of Human Centered Computing Technical Area), "Human-Centered Computing: Methods, Program, and Progress", March 26, 1999

Dr. Clancey offers an important approach to modeling and improving human productivity within organizations. Dr. Michael Shafto, Chief of Human-Automation Integration Research Branch, will introduce Dr. Clancey.

In this presentation I will show how human-centered computing has developed at Ames into a research program with a portfolio of new projects and cross-center partnerships. I will explain the key ideas of participatory design and the total systems perspective that frame HCC research, and

contrast HCC with traditional human factors and software engineering. I will describe how we are using the Intelligent Systems program to set up cross-center partnerships and outline a "project maturity model" by which progress over the next five years may be expected. Examples will show how HCC research covers the gamut from facilitation in computer system projects to inventing new modeling languages.

Dr. Michael Shafto (Human-Automation Integration Research Branch); Dr. Erik Vinkhuyzen (RIACS), "Communication and Coordination in Mission Control", March 19, 1999.

Dr. Vinkhuyzen will show a video gathered during his work with the Flight Dynamics Officers in the Mission Control Center at JSC. The analysis of the materials will focus on the singular communications technology used by the Flight Controllers in Mission Control, the "voice loop system." Of special interest is the way individuals have adapted their natural communication skills to the capabilities and limitations of the voice loop system.

Dr. Steven Zornetzer (Drector of Information Systems Office), "Overview of IT Research at ARC: Challenges and Opportunities", March 05, 1999

Dr. Zornetzer's talk will have broad appeal because NASA Administrator Dan Goldin has designated the Ames Research Center as NASA's Center for Excellence in Information Technology. Most RIACS scientists work in Code I, for this reason also the talk will be of special interest for us at RIACS. This will be a good time to meet Dr. Zornetzer, learn more about his vision for information technology and ask questions.

Dr. Muriel Ross (Code SL); Dr. Marjory Johnson (RIACS); Dr. Alexander Twombly (RIACS), "Center for Bioinformatics and the Joint BC/NREN Telemedicine Experiment", February 26, 1999

The Center for Bioinformatics is dedicated to developing advanced, high fidelity 3-D imaging and interactive virtual environment technologies for biomedical and scientific purposes. This talk will emphasize those efforts directed toward the biomedical community. 3-D imaging is carried out from serial sections of biological and medical tissues and organs, whether these are obtained from microscopy or from CT or MRI scans. A special effort is underway to produce similarly high fidelity images from echo and sonic data, since these are possible to obtain on space station and spacecraft, or on distant planets. The images are to be used for diagnostic purposes and to simulate, in virtual environment, surgical procedures for planning and training purposes in medical settings on Earth. The long-term goal is to ensure the health of astronauts as they probe deeper into space, so that these same technologies will eventually find use on spacecraft should unanticipated medical problems arise. A natural outcome of this research is telemedicine, in which the goal is to bring the clinic to the patient rather than the patient to the clinic. Communication to

spacecraft for medical advice is, after all, but another example of reaching into a remote site via telecommunication.

Dr. Sean Casey, Sr. Scientist (SOFIA Project), "SOFIA: NASA's Stratospheric Observatory for Infrared Astronomy", February 19, 1999.

I plan to present an overview of the facility covering development interests for both the observatory and science instrument programs. The question and answer period should cover the application of information research concepts to SOFIA as a demonstration platform."

Dr. Peter Norvig, Chief of Computation Sciences Division (IC); Dr. Peter Cheeseman, (RIACS), "Integration of information from satellite images", February 12, 1999.

Dr. Norvig will outline IC Division projects. Dr. Cheeseman will discuss the integration of information from satellite images.

Dr. David Morrison, Director of the NASA Ames Space Directorate, "Astrob Febraury 05, 1999.

Dr. Morrison's directorate embraces a literal universe of scientific areas of intellectual interest, including (just a sampler): origin and evolution of galaxies, stars and planetary systems, origin and evolution of life, and possibilities for humans to inhabit other worlds. Some of Dr. Morrison's enabling goals are to establish a virtual presence throughout the solar system and develop revolutionary technologies for missions.

IV   RIACS Staff

The primary mechanism used at RIACS is the engagement of RIACS scientists in NASA projects. RIACS staff scientists are recruited from the broad university research community to provide a nucleus of activity to both collaborate with NASA scientists and to provide an "attraction point" for bringing in visitors from academia. These scientists typically work in clusters on NASA projects – collaborating closely with NASA scientists and onsite contractors, providing senior technical expertise and coordinating visitors from academia.
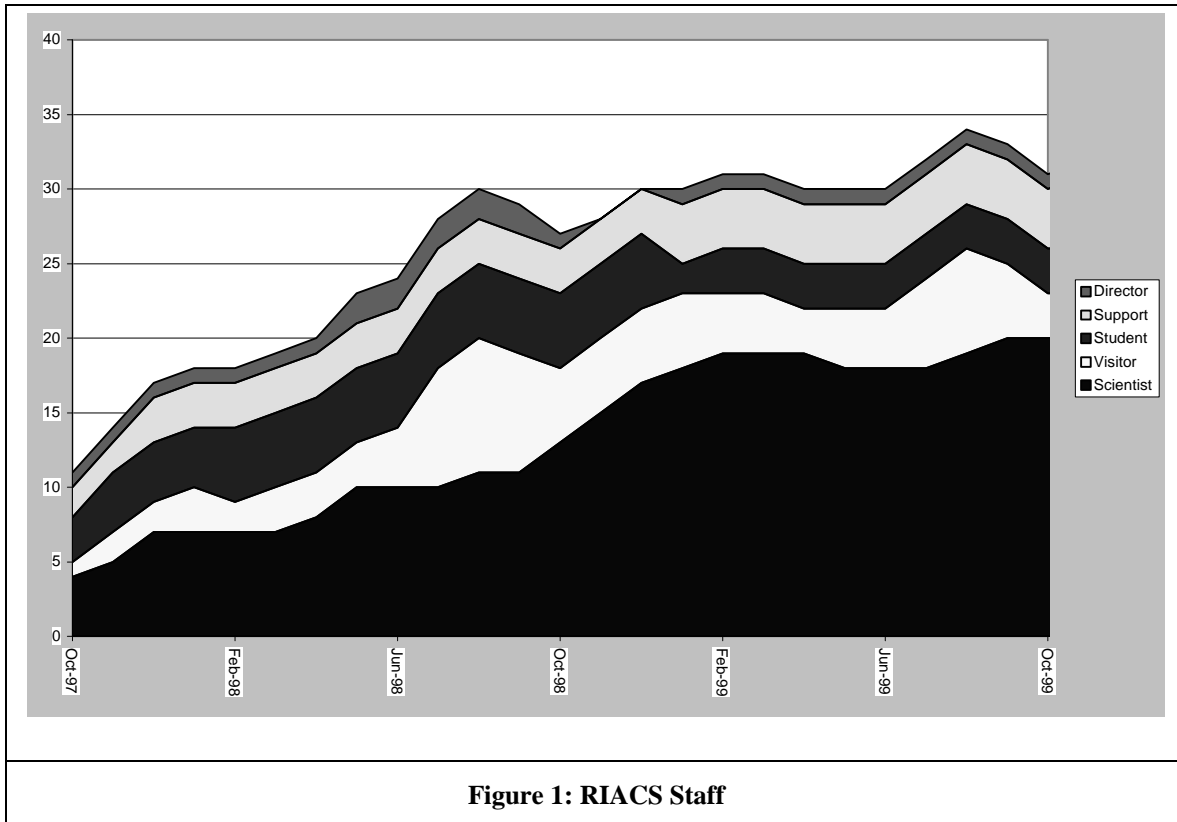

Figure 1: RIACS Staff

Figure 1 shows the growth in RIACS staff from the beginning of 1997 through the reporting period. Over the period of the current Cooperative Agreement shown, the number of scientific staff onsite has grown from three to twenty.

The following sections lists both the permanent and visiting staff at RIACS over the reporting period, with the time they have been at RIACS shown in parentheses.

IV.A   Management Staff

There were two changes in the Director during the year. Dr. Robert Moore left RIACS at the beginning of the year. Dr. Michael Raugh served as Interim Director while a search for a new Director was conducted. In June 1999, Dr. Barry Leiner was appointed Director. In July 1999, Dr. Robert Morris joined RIACS as Deputy Director.

Barry Leiner, Director – Ph.D. Electrical Engineering, Stanford University, 1973. (06/28/99 – present)

Michael Raugh, Interim Director – Ph.D. Mathematics, Stanford University, 1979. (12/28/88 – 06/28/99)

Robert C. Moore, Director – Ph.D, Artificial Intelligence, Massachusetts Institute of Technology, 1979. (04/13/98 – 10/31/98)

Robert Morris, Deputy Director – Ph.D, Philosophy, Indiana Universitiy, 1984. (08/02/99 – present)

IV.B   Administative Staff

Diana Martinez, Administrator (08/18/97 - present)

Beatrice Burnett, Administrative Assistant (11/3/97 - present)

Rasheeda Shaheed, Administrative Assistant (04/16/99 – present)

Daidra Gibson, Administrative Assistant (12/28/98 – 04/19/99)

Steven Suhr, Systems Administrator (11/1/96 – 09/17/99)

IV.C   Scientific Staff

Esfandiar Bandari, Ph.D, 1995, Computer Science, University of British Columbia, Computational Vision, Signal Processing, 3-D Reconstruction and Medical Imaging (09/2/98 – present)

Peter J. Cheeseman, Ph.D., 1979, Monash University, Artificial Intelligence, computational complexity, baysian inference, computer vision, plasma physics (09/1/97-present)

Rei J. Cheng, BS, 1973, Pharmacy, Taipei Medical College, 3-D scientific visualization and medical visualization (09/1/98 – present)

Bernd Fischer, MS., 1990, TU Braunschweig, Germany, Computer Science (11/01/98 – present)

Dave Gehrt, JD Law, University of Washington, 1972, UNIX system administration, security, and network based tools (1/84 - 7/85, 2/1/88 - present).

Beth Ann Hockey, Ph.D., 1998, Linguistics, University of Pennsylvania (01/01/99 – present)

Frances H. James, Ph.D., 1998, Computer Science, Stanford University, Speech understanding technologies and applications, developing interfaces using speech.
(7/13/98 – present)

Marjory J. Johnson, Ph.D., 1970, Mathematics, University of Iowa, High-performance networking for both space and ground applications (1/9/84 -present).

Ari K. Jonsson, Ph.D., 1997, Computer Science, Stanford University, New Techniques in constraint satisfaction and Automated Planning. (10/1/97 - present)

John Loch, Ph.D., 1998, Computer Science, University of Colorado at Boulder (12/28/98 – present)

Robin D. Morris, Ph.D., 1995, Engineering, Trinity College, Cambridge, U.K., Bayesian Inference and Image Analysis, Mathematical Modeling of Natural Imager and Associated Computional Methods. (3/1/98 – present)

Pandurang Nayak, Ph.D., 1992, Computer Science, Stanford University, 1992, Autonomous systems, model-based reasoning, diagnosis and recovery, qualitative and causal reasoning, abstraction and approximation in knowledge representation and reasoning. (11/15/97 – 10/1/99)

John O'Neill, Ph.D., 1997, Computing, Royal Melbourne Institute of Technology (12/01/98 – present)

Seungjoon Park, Ph.D., 1996, Electrical Engineering, Computer Science, Stanford University (10/01/98 – present)

Lewis Peach (09/04/99 – present)

Charles Pecheur, Ph.D., 1996, Electrical Engineering, Computer Science, University of Liege, (11/06/98 – present)

Barney Pell, Ph.D., 1993, Computer Science, Cambridge University, Autonomous Agent Operations of control robotic and software control systems. (11/1/97 –12/31/98)

Manny Rayner, Ph.D., 1993, Computer Science, Royal Institute of Technology, Stockholm, (07/01/99 – present)

Maarten Sierhuis, MS, 1986, Engineering, Mague Polytechnic University, Intelligent multi-agent simulation agent-oriented programming languages (4/1/98 – present)

Ian A. Twombly, Ph.D., 1997, Biophysics, John Hopkins University, 1997, (9/1/98 – present)

R. Erik Vinkhuyzen, Ph.D, 1997, Psychology, Zurich University, study human aspects of technology use (4/6/98 – 4/29/99)

Willem Visser, Ph.D, 1998, Computer Science, University of Manchester (10/19/98 – present)

IV.D Visiting Scientists and Consultants

(Note: Several visitors had multiple stays at RIACS. Dates shown are inclusive)

Remi Abgrall, Ph.D. – Professor, University of Bordeux, France, Numerical implementation of boundary conditions for first order Hamilton Jacobi equations. (7/27/98 – 8/30/99)

Turan Coratekin – Lehr und Forschungsgebiet Fuer Mechanik, Turbulence modeling and computational fluid dynamics. (6/13/99 – 12/12/99)

Ronald Henderson, Ph.D. - Sr. Scientist, California Institute of Technology, Computational Fluid Dynamics, parallel computing, hydrodynamic stability, turbulence (2/17/98-7/23/99)

Richard G. Johnson, Ph.D. - Physics, Indiana University, 1956, Global environmental problems and issues (11/1/92 - present).

Anil Kokaram, Ph.D. – Lecturer, University of Dublin, Trinity College, Ireland, Video processing and image dquences. (7/30/99 – 9/3/99)

Dimitrios Maroudas, - Asst. Professor, Chemical Engineering, U.C. Santa Barbara, Theoretical and Computational Materials Science with emphasis on surface science and microstructure evolution in semiconductors, metallic thin films and structural alloys (9/22/97-9/4/99)

Bernd Muller, Ph.D. – Lecturer, Upsala University, Sweden, Turbulence simulation. (7/12/99 – 7/23/99)

Neil Sandam, Ph.D. - Lecturer, Queen Mary & Westfield College, England, Direct numerical simulation of transitional and turbulent fluid flow, turbulence modeling (12/30/97 – 8/25/99).

Steven Senger, Ph.D. – Professor, University of Wisconsin, LaCrosse, Filtering and visualization of CT and Ultrasound data sets (9/1/98 to 7/31/99)

Andrew Sohn, Ph.D. - Assistant Professor, New Jersey Institute of Technology, Parallel computing, high performance computing, distributed-memory with multi-processor architectures, dynamic load balancing for adaptive computations fast mesh partitioning (6/13/98-1/15/99).

Sankaran Venkateswaran, Ph.D. -Sr. Research Associate, University of Tennessee Space Institute, preconditioning methods and use of neural networks for application to semi-conductor materials (5/18/98 – present)

David Zingg, Ph.D. - Associate Professor, University of Toronto, Canada, Development and analysis of high-accuracy numerical methods applicable to simulations of fluid flows, acoustic waves and electromagnetic waves (6/21/99-8/27/99).

IV.E   Visiting Students

Amrita Bhagat – Computer programming, DeAnza College. (7/26/99 – present)

Susie Go, Ph.D. – 1998, Applied Math, University of California, Los Angeles, multilevel methods on unstructured grids. (1/21/96 – 12/31/98).

Roy Lue – Internet Based TCAD using IPG, Stanford Unversity, Stanford, California, Center for Integrated Systems. (1/1/98 – 12/31/98)

Tae-young Oh – Device modeling and simulation, Stanford University. (1/1/98 – present)

Leonid Oliker, Ph.D. - University of Colorado, compilation of data parallel programs. (9/1/94 – 12/31/98).

Shyam Ramalingam - Chemical Engineering, University of Santa Barbara, Computer Modeling of desposition of Si Films using PECVD. (9/2/97-present)