

COMPUTER SCIENCE

RESEARCH PROJECTS 2000

CAREER: Towards Estimating Requirement Coverage: Managing Goals and Scenarios during Requirements Evolution

Ana I. Anton

\$220,000 by the National Science Foundation

June 2000 - May 2004

This proposal focuses on efforts to integrate core research and educational objectives. The research addresses important issues in the discovery, elaboration and management of system use scenarios for the specification of software requirements. The ultimate goal is to develop viable solutions for supporting the early stages of the software lifecycle by ensuring requirements coverage.

Toward Object-Oriented Rapid Prototyping and Automated Code Generation

Ana I. Anton

\$75,108 by Bell South

May 2000 - May 2001

The objective of this research is to compare and contrast two object-oriented software development approaches software approaches via elaboration and software development via translation. Our investigations will compare the pros and cons of these methods and all findings will be demonstrated through non-trivial examples involving telephony features and services. Additionally, we expect to be able to show how the adoption of certain tools and approaches can lead to more rapid prototyping and automatic code generation.

Creating and Validating CMM-Compatible Software Processes Models

Ana I. Anton

\$34,378 by Asea Brow Boveri (ABB)

August 2000 - June 2001

The objective of this research is to create and validate CMM (Capability Maturity Model) compatible software process models for adoption within the software development groups at ABB and the NCSU electronic commerce project studio. Investigations will

compare various process models so that they may be tailored to meet different organizational needs.

A Perceptual Visualization Architecture

Christopher Healey

\$354,000 by the National Science Foundation

September 2000 - August 2003

This project will address three important issues: (1) can we harness and apply low-level human perception to the problem of visualizing large, complex, multidimensional datasets? (2) can we embed this knowledge in an AI-based system that will assist viewers in constructing perceptually-optimal visualizations, and can we do this in a general way to address a wide range of problem environments? and (3) can results from perception be bound to stylistic properties in Impressionist painting, thereby creating a system that allows a viewer to "paint" an expressive visual representation of their data, while at the same time ensuring the result accurately portrays the underlying data values being displayed?

Interactive Exploration of Complex Datasets Via the Effective Generation of Text and Graphics

Christopher G. Healey, Robert St. Amant, R. Michael Young

\$569,338 by the National Science Foundation (IIS/IDM)

September 2000 - September 2003

This is accomplished through an interlocking collection of textual and graphical discourse engines, a plan recognition system, and an interaction manager that harnesses the power of these tools. Users begin by asking questions about the data stored in their dataset. The system answers using a combination of text and graphics. Text responses are built by a discourse engine, while graphical images are constructed using a perceptual visualization assistant. The collection of all possible responses is evaluated to select the most effective answer, be it text, graphics, or a combination of the two. A plan recognition system is used to analyze the users' queries and their reactions to the responses they receive. This allows the system to anticipate future queries, cache relevant statistics, and guide the discourse and visualization systems during the evaluation and selection of appropriate answers to each user query. Results from this project will include: (1) plan recognition and interaction plan construction performed by the system to identify and model current and future analyses conducted by the users; (2) presentations that are sensitive to both the current and anticipated future state of users' investigations; (3) assisted navigation techniques; (4) methods for evaluating the effectiveness of the use of text and/or graphics; and (5) perceptual visualization techniques. Results will be disseminated through journal and conference publications, online datasets with results and analysis, and online software demos of relevant research components. Although this project will study applications from the oceanography and public school domains, these results are relevant in any situation where interactive exploration of large, complex datasets is required.

CAREER: Assisted Navigation in Large Visualization Spaces

Christopher G. Healey

\$370,403 by the National Science Foundation (ACIR/ACR)

February 2001 - January 2006

This project is an investigation of methods for assisting with the navigation of large, complex information spaces. Although results from these studies are relevant to a number of research areas, interest will be focused on construction of a navigation system designed to help viewers visualize, explore, and analyze large, multidimensional datasets. Methods to assist with the analysis and navigation of these types of datasets was specifically cited as an important open problem by the joint DOE/NSF panel on future research in visualization. The work will combine a detailed local display and a high-level global overview of the locations and structure of areas of interest within the dataset. The local view will use perceptual cues to harness the abilities of the low-level human visual system. The global overview will be built in two separate stages. First, elements of interest will be identified using a combination of: (1) explicit rules provided by the viewer, and (2) implicit rules built by watching what viewers select, where they move, and what they examine. Next, the elements will be clustered into one or more areas of interest. The use of graph construction techniques like planar triangulations and minimum spanning trees will be investigated to link the elements together. An underlying graph that: (1) supports efficient navigation via the application of graph traversal algorithms, and (2) provides an effective global overview to visualize the areas of interest and the relationships that exist between them will be sought. A set of validation experiments will be designed to identify the strengths and limitations of our navigation techniques. Datasets from the oceanography and e-commerce domains will be used to test the system in a practical, real-world environment. The first set of experiments will work with domain experts, in part to provide anecdotal feedback on our system, and in part to identify fundamental navigation and exploration tasks performed during visualization. These tasks will then be integrated into a controlled experiment that studies the performance of our system vis-a-vis a system without navigation aids, and existing focus+context visualization techniques specifically designed to display these types of large, complex datasets.

Triangle Computer Science Distinguished Lecturer Program

S. Purushothaman Iyer

\$36,000 by the U.S. Army

August 1998 - July 2001

The Triangle Computer Science Distinguished Lecturer Series was started in 1995 based on a grant from ARO and is jointly administered by the Computer Science departments at Duke University, North Carolina State University and UNC Chapel Hill. The original motivation for the lecture series was to educate and inform the students at the three universities and local computer professionals about exciting new research being

conducted across all areas of computer science. In the second phase of the series (during the years 1998-2001) we propose to continue with the momentum in building a truly superb educational experience for our students and faculty, and to computer professionals in the Triangle area.

Automatica Analysis of Probabilistic Systems

S. Purushothaman Iyer

\$258,648 by the U.S. Army Research Office

July 2001 - June 2004

Concurrent systems such as network protocols and net-centric programs are difficult to build and debug because of the potential they exhibit for unintended process interactions. The development of net-based applications which have to contend with probabilistic guarantees from lower-levels is even more difficult as they need to be functionally correct and also satisfy reliability/performance constraints. This project will investigate how formal methods can be extended to address both logical correctness and reliability/performance constraints. The current project will explore semantic theories of systems that have both non-determinism and probabilistic choice. In particular, notions of equality and approximate equality of system behaviors will be investigated. Furthermore, the effect of these notions on compositional reasoning will also be studied. The second topic of the proposed work will be a thorough comparison of the semantic theories developed in this project against traditional approaches to dealing with non-determinism and probabilistic choice. Finally, practical algorithms for process minimization and for checking equality (and approximate equality) of processes will be designed and implemented in the Concurrency Workbench of New Century. Case studies, to evaluate the proposed theories, will also be constructed and studied.

Abstraction-Based Approaches to Correct Reactive Software

S. Purushothaman Iyer, Rance Cleveland

\$270,000 by the U.S. Army

July 1998 - June 2001

It is by now a well accepted fact that checking designs for errors early in the software development cycle could lead to early identification of errors, and their correction. However, once designs have been verified there is no guarantee that software implementations implement what has been designed. In this project, we have been investigating how abstract interpretation techniques (developed in the context of compilers) can be used to abstract programs and how the ensuing abstractions can be checked against designs. The technical questions we have been dealing with are: (a) how reasonable approximations can be derived from distributed programs, and (b) what approximation techniques guarantee that results about abstractions hold of original programs.

A Data Mining Approach for Building Cost-Sensitive and Light Intrusion Detection Models

Wenke Lee

\$2,001,561 by DARPA; Subcontractors: Columbia University and Florida Institute of Technology.

August 2000 - August 2003

This project proposes a novel system for rapid development and deployment of effective and cost-sensitive IDSs. Our system automates feature construction, the critical step in building effective misuse and anomaly detection models, by analyzing the patterns of normal and intrusion activities computed from audit data. Detection models are constructed automatically using cost-sensitive machine learning algorithms to achieve optimal performance on the given (often site-specific) cost metrics. Our system finds the cluster of attack signatures and normal profiles and accordingly constructs one light model for each cluster to maximize the utility of each model. A dynamically configurable group of such light models can be very effective and efficient, and resilient to IDS-related attacks. The detection models can be very rapidly deployed through automatic conversion to efficient real-time modules of fielded IDSs.

Integration of Intrusion Detection and Network Management: A Feasibility Study

Wenke Lee

\$20,000 by NCNI Fellowship with funding from Aprisma Management Technologies (second year in a row)

July 2000 - June 2001

It provides support for a graduate student for one year and the project will study whether and how information gathered by a to improve the detection accuracy. The main tasks of the project include network management system can be utilized by an intrusion detection system developing algorithms for anomaly detection and early sensing of intrusions.

Self Explaining Learning Environments

James Lester, Patrick Fitzgerald

\$962,027 by the National Science Foundation

August 1999 - July 2002

Given the centrality of explanation in science, one of the most intriguing forms of intelligent multimedia learning environments revolves around explanation. Self-explaining learning environments will provide two complementary explanation functionalities: immersive explanation generation, which will enable learners to pose questions to objects and processes in 3D worlds that can dynamically generate clear cinematic and narrative explanations of their own structure, function, and causality. Second, student-driven explanation construction systems will pose problems about phenomena under investigation by students. To solve the problems, students will employ

a graphical explanation construction system to build explanation structures that will be semantically analyzed. If successful, the net result of this work will be two-fold: (1) a suite of intelligent, learning environment technologies whose communication abilities and, hence, whose learning effectiveness, significantly surpass current-generation educational software, and (2) a comprehensive theory of explanation-based constructivist learning that furnish a data-rich set of principles for guiding the design of learning environments.

Animated Pedagogical Agents for Constructivist Learning Environments

James Lester, Sharolyn Converse, Patrick Fitzgerald, Hiller Spires
\$600,472 by the National Science Foundation
September 1997 - August 2001

Constructivist learning environments facilitate discovery-based learning through creative problem-solving experiences. However, precisely the characteristics of constructivist learning which offer the greatest potential also pose the greatest challenge: the complexity of the learning experience, which follows from the freedom to engage in exploratory learning activities such as designing artifacts in microworlds, constantly threatens to overwhelm learners. To combat this complexity, learning environments should provide scaffolding in the form of highly contextualized problem-solving advice that is customized to each learner. Perhaps the most intriguing vehicle for providing such dynamically individualized scaffolding is the technology of animated pedagogical agents. Featured prominently in learning environments, they couple key feedback functionalities with a strong visual lifelike presence to support effective constructivist learning. In addition to the cognitive implications of animated pedagogical agents' ability to provide visually contextualized advice, the immediate and deep affinity that children seem to develop for them suggests that the direct pedagogical benefits they offer are perhaps even exceeded by their motivational benefits. Imbued with the illusion of life, dynamically animated agents have the potential to significantly increase the time children seek to spend with learning environments. The twin objectives of imated pedagogical agents, and (2) provide a comprehensive account of the cognitive processes this research are: (1) create a new generation of intelligent constructivist environments with an and results of interacting with constructivist learning environments with animated pedagogical agents.

CAREER: Multimedia Explanation Generation for Knowledge Based Learning Environments

James Lester
\$390,374 by the National Science Foundation
June 1997 - May 2001

Knowledge-based learning environments can provide highly customized problem-solving experiences that are tailored to the individual needs of each student. Explanation generation is the key functionality of learning environments that are designed to provide knowledgeable feedback. As a result of rapid advances in multimedia technologies, we are now presented with the opportunity to create real-time multimedia explanation

generators that employ animated pedagogical agents. The primary objective of this research is to create the technology for a new generation of knowledge-based learning environments. It has three major thrusts: (1) developing a computational model of real-time multimedia explanation planning that can be employed to construct pedagogical customized, multimedia explanation plans, (2) developing a computational model of animated pedagogical agents that actively participate in a multimedia explanation, and (3) conducting formal empirical evaluations of the pedagogical effectiveness of the model of explanation generation.

Visual Factors Affecting Pilot's Judgments of Trajectory-to-Touchdown During Emergency Landings

David McAllister, Donald Mershon, Celeste Mayer

\$185,999 by the National Aeronautics & Space Administration

January 2000 - December 2001

The research will employ the VisionDome® maintained at the School of Design, NC State University, to simulate emergency landings fields with various amounts of structure (e.g., trees, buildings) around the periphery. Participants will be seated within a structural mockup of a typical single-engine airplane cockpit. Cockpit instruments will be visible, but not functional. At times, the representation of a "windmilling" propeller will also be visible, at an appropriate distance in front of the participant. Participants will be passive observers in an airplane approaching an emergency landing field. Participants will be randomly presented with a series of trials consisting of all combinations of different constant-bank-angle approaches and three different landing trajectories which would result in a landing prior to an optimal touchdown point, at the touchdown point or beyond the touchdown point. Participants, both experience pilots and naïve observers, will be asked to indicate their expected touchdown point at specific times during each approach, as well as their perceived bank angle. Analyses of the passive-observer trials will be used to guide further investigation. This will enable later attention/effort directed toward those situational conditions most likely to lead to significant judgmental errors.

Fault Tolerant Networking through Intrusion Identification and Secure Compartment

Douglas Reeves

\$394,065 by Defense Advanced Research Projects Agency

July 2000 - July 2003

We will study the concept of Compartment in a networking system to handle different types of network infrastructure attacks. We particularly will develop a compartment-based network system to protect BGP (Border Gateway Protocol) routing protocol. Furthermore, we will combine Compartments with the Deciduous model to more effectively identify the intrusion sources.

Protecting Network QoS Against Denial of Service Attacks

Douglas Reeves, Peter Wurman

\$1,448,788 by the U.S. Air Force

August 1999 - June 2002

The next generation Internet will have mechanisms for providing Quality of Service (QoS). All the existing QoS developments are concentrating on how to make QoS guarantees/assurance in the absence of hostile elements. There are a few instances in which mechanisms for protecting the integrity of reservation messages have been considered. In this project, we address three new vulnerability questions with regard to QoS provisioning: (1) how to prevent "Killer Reservation" denial-of-service (DoS) attacks; (2) how to detect and respond to DoS attacks that are directed at the actual data-flow; (3) to what extent resource reservation and allocation mechanisms can be designed to reduce their vulnerability and to respond effectively to various DoS attacks. All these questions will be explored with the objective of improving the inherent survivability of the QoS mechanisms. We recognize that DoS attacks are difficult to deal with in general. Our approach to this problem consists of: (1) applying a new mechanism called congestion pricing in resource reservation and allocation process to protect against vicious end-customers; this mechanism provides policy-based control over network resource utilization, and is simple to administer and monitor; (2) using an intrusion detection and response component together with the QoS mechanisms; this component provides detection of DoS attacks on the data-flow, and isolates the attacking packet streams by dynamically remarking selected packet streams in coordination with other bandwidth brokers; (3) enhancing the security of the resource reservation and setup protocol (e.g., RSVP) against insider router attacks.

A Competitive Market Approach to Distributed Resource Allocation with QoS and Priorities

Douglas Reeves

\$141,059

May 1999 - November 2001

The quality of service (QoS) of an application is determined by the amount of resources it needs, and the amount of resources which are available and allocated to it. Current methods of resource allocation are either very conservative (based on static worst-case assumptions), or they fail to recognize the different requirements of different applications (resulting in equal allocation to everyone). We are developing a method of resource allocation which is dynamic, distributed, provides equal QoS rather than equal resource amount, and which can also consider the relative importance or priority of different applications. This method is based on competitive pricing of resources in response to demand, and the available budgets of users. This method has been applied to flow control in ATM networks. We propose to apply it also to server allocation, congestion control in the Internet, and access scheduling in wireless networks. We propose to extend our

theory to permit a limited form of reservations and admission control. By interpreting our prices as real money, we expect to be able to justify economic decisions about which resource provider to choose, when to expand the resource capacity, and how to generate revenue. For networks, we also propose to apply our results in resource allocation for the control of delay bounds for real-time traffic, and we plan to investigate the issue of resource allocation for multiple receivers of a single traffic stream (multicasting).

Flow Control for Multimedia Streaming Over the Internet

Injong Rhee

\$300,000 by the National Science Foundation

June 2000 - May 2002

As the Internet becomes more diversified in its capabilities, it becomes more feasible to offer services that were not possible under earlier generations of Internet technologies. Real-time multimedia streaming and IP multicast are two such emerging technologies. The objective of this proposed work is to develop, verify analytically and experimentally, and implement a suite of end-to-end flow control protocols for unicast and multicast real-time streaming applications. The developed protocols are evaluated based on fairness, TCP-friendliness, stability and scalability. These properties must hold regardless of the types of networks, or more specifically whether networks are symmetric or asymmetric in bandwidth and delays.

Congestion Control for Unicast and Multicast Streaming Application

Injong Rhee

\$22,933 by the Center for Advanced Computing and Communications

July 2000 - June 2001

This proposal is a follow up on the previous CACC projects on congestion control for reliable multicast. The following are issues that have not been adequately addressed in the previous projects: (1) Unicast: we need to evaluate its performance under more realistic Internet environments. In particular, we are interested in its performance under a large number of on and off traffic flows. (2) Multicast: we plan to develop flow control for receiver-driven layered multicast (RLM). RLM ensures scalability and inter-receiver fairness by allowing receivers to determine their own receiving rates independently. However, its stability and TCP friendliness have not been proven. Using TEAR, we will investigate ways to improve these properties under heterogeneous network environments. (3) For multicast data transmission that is not amenable to layering, we plan to develop flow control techniques for sender-driven single rate multicast (SSRM).

CAREER: Investigation of Error Recovery Techniques for Interactive Video Transmission over Wireless Networks

Injong Rhee

\$269,075 by the National Science Foundation

April 1999 - March 2003

We propose to study a new class of error recovery techniques that focuses on eliminating error propagation. Prior work has mainly taken an approach that prevents display errors from occurring from the first place. This existing approach is not effective for wireless interactive video transmission because data losses inevitably occur in wireless communication, and repairing them tends to introduce the delays in video playout. Our approach is to isolate errors when they occur by preventing them from propagating. The delays in repairing data losses affect only the duration of error propagation. We call our approach Recovery from Error Spread using Continuous Updates (RESCU). RESCU does not introduce any delay in video playout, and has potential to achieve good error resilience. A similar approach is adopted in several recovery techniques stipulated in H.263+, an International Telecommunication Union (ITU) low-bit rate standard. However, they are limited in that they require feedback channels and are not suitable for multicast. Our proposed techniques, in contrast, have potential to work well with or without feedback channels, and to be scalable for multicast. Encouraging preliminary results indicate that such potential is highly realizable.

CAREER: Safe Observable Components: Programming-Language Support for Object Interaction Protocols

Jonathan Rossie, Jr.

\$107,00 by the National Science Foundation

July 1999 - June 2001

The research focuses on programming-language support for the behavioral characterization of observable components in terms of logical finite-state machines. Observable components support two-way messaging protocols with their client code, keeping the clients informed of significant state changes that affect the component's outward utility. The logical state machine provide a vocabulary for this protocol, serving as enforceable characterizations of the components' dynamic protocol. The research objective is to develop practical language mechanisms, rigorous formal foundations, and efficient implementation techniques to support observable components and to exploit the enforceable state machines to improve the safety of component reuse. A safe implementation would protect components from receiving messages out of sequence, and would also protect clients from components that are not faithful to their declared state machines. An essential technology in developing and using these components is a notion of safe substitutability that would account for the state machine and its implied two-way protocols. The research plan includes formal foundations for safe substitutability of components, as well as robust and efficient implementation technology in sequential and concurrent contexts.

Regional Testbed Optical Access Network for Internet Protocol (IP) Multicast and Differentiated Services (Helios)

George Rouskas

\$224,825 by Microelectronics Center of North Carolina

March, 2000 - March 2002

The proposed research will develop a regional optical testbed network carrying IP traffic between MCNC, UNC, NCSU, and Duke. The testbed architecture will consist of interconnected access networks. The access networks and the interconnecting network will utilize an optical single-hop architecture. Network elements consist of access nodes interconnected to a passive coupler to form a broadcast star topology. Access node implementations will utilize an extended version of the HiPer-1 optical Media Access Control protocol developed and analyzed at NCSU.

Multicast Support for Just In Time Optical Burst Switching Networks

George Rouskas, Harry Perros

\$69,642 by Microelectronics Center of North Carolina

October 2000 - March 2001

Researchers will develop and demonstrate technologies to support multicasting for Just in Time signaling protocols for optical burst switching networks. We assume that WDM switches support multicasting at the optical layer. In other words, the switches have the ability to (1) optically split a signal arriving on any wavelength at an input port to create two or more identical copies of the signal, (2) independently switch each copy of the signal to a different output port, (3) amplify each copy of the signal to prevent degradation, and (4) (possibly) convert each copy of the signal to a different wavelength for transmission out of the switch. We will extend the signaling protocol to provide a best-effort setup for uni-directional point-to-multipoint data transmission.

Performance Analysis of Wavelength Routed Wide Area Networks

George N. Rouskas, Harry G. Perros

\$255,384 by the National Science Foundation, Division of Advanced Networking and Infrastructure Research

September 1998 - August 2001

Wavelength routed wide area networks will be an integral part of the future all-optical network infrastructure. It is envisioned that these networks will act as the backbone that provides the interconnection for local area lightwave sub-networks attached to them. A significant amount of research effort has been devoted to the design and evaluation of wavelength routed architectures in recent years. However, the deployment of an

infrastructure that will enable large scale interconnection of optical networks raises a number of fundamental and challenging problems that require novel and innovative solutions. The research plan outlined in this proposal identifies a number of significant but open problems arising in wavelength routed environments, and presents a comprehensive and novel approach to address them. The outcome of the proposed research will be an analytical framework that will permit us to efficiently compute call blocking probabilities in wavelength routed networks, evaluate wavelength assignment policies and converter placement strategies, investigate routing algorithms for arbitrary network topologies, and study the performance of multicast calls. One of the main objectives of our work is to identify potential fairness problems that may arise in such an environment, and to employ a combination of routing, converter placement, and wavelength assignment strategies to alleviate these problems and ensure fair treatment for all calls.

CAREER: Towards an All-Optical Network Infrastructure: Interconnection of Photonic WDM Broadcast-and-Select Local Area Networks

George N. Rouskas

***\$200,000 by the National Science Foundation CAREER Program, Division
of Networking Research
September 1998 - August 2001***

The career development plan outlined in this proposal focuses on the integration of the research and teaching objectives of the principal investigator. The research project addresses a number of important issues in the design and interconnection of optical networks. The ultimate goal of the proposed work is to develop viable solutions for a seamless nationwide all-optical network infrastructure. The educational component includes development of graduate-level courses on all-optical networking and on internetwork architectures and protocols, integration of topics on lightwave technology in existing graduate and undergraduate introductory network courses, and active involvement of undergraduate students in network projects. In an all-optical network using wavelength division multiplexing (WDM), information is transmitted as optical signals throughout the path from source to destination. By using multiple concurrent channels, each on a different wavelength, WDM has the potential to exploit the enormous information-carrying capacity of fiber and deliver an aggregate throughput that is several orders of magnitude higher than it is possible today. Given the installed base of optical fiber, and the maturing of optical component technologies, it appears that current network technologies will eventually evolve to an all-optical infrastructure. The deployment of such an infrastructure will cater to the needs of emerging bandwidth-intensive applications. However, large scale interconnection of optical networks raises a number of fundamental and challenging problems that require novel and innovative solutions. As part of his career development plan, the principal investigator proposes to address a number of these issues, including multiple access in broadcast-and-select WDM networks, traffic scheduling of unicast and multicast traffic, interconnection of photonic local area networks, architectural issues in wavelength routed backbone networks,

distributed control, and issues related to the interaction among subnetwork and backbone protocols. Our work is expected to further our understanding of the characteristics of optical networks, and to provide insight on how to overcome some of the limitations of the technology. The objective is to produce an analytical framework that will make it possible to evaluate alternative architectures and protocols.

Combinatorial Generation, Enumeration, and Structure.

Carla D. Savage

\$20,000.00 by the National Security Agency (DOD)

February 2000 – August 2001

The proposed research is an investigation of fundamental questions involving the structure and size of combinatorial families and various methods for counting and generating them. Efficient algorithms and combinatorial structure go hand in hand. Insight into structure leads to more efficient algorithms which, in turn, allow a larger view of families whose size can grow exponentially.

Multiagent Systems for Networking and Telecommunications

Munindar P. Singh

\$68,000 by Cisco

January 2001 - December 2001

The rapid and unprecedented expansion of the networking infrastructure available bandwidth is pushing the decentralization of control across the various layers of distributed computing architectures. There is more intelligence at the fringes of the network than ever before; there is also greater intelligence sprinkled throughout the network. This research project will study the ramifications of this paradigm shift in networking and telecommunications. A specific problem being studied is adaptive packet routing.

Multiagent Referral Systems for Knowledge Management

Munindar P. Singh

\$40,000 by IBM Research

August 2000 - August 2001

Knowledge management is a major challenge facing users worldwide, especially within enterprises of all sizes. This project develops an approach based on multiagent systems that will enable users to find the most relevant sources of knowledge within their enterprise.

Computational Principles of Trust

Munindar P. Singh

\$476,574 by the National Science Foundation

September 2000 - August 2003

Successful interaction relies heavily upon trust. This applies equally to electronic commerce and virtual social communities. However, figuring out who to trust and to what extent is extremely difficult in open networked information environments. Trust is a complex concept and involves aspects of competence and good nature (of the trusted party) and the risk tolerance and urgency (of the trusting party). This project will study distributed, scalable computational approaches for trust management taking into account the different forms in which trust is exhibited in networked environments. This project studies aggregate phenomena such as the emergence of subcommunities, pivots (which link different subcommunities), and the sensitivity of a community to the addition or removal of a member. CAREER: Cooperative Information Systems Munindar Singh \$261,185 by the National Science Foundation August 1996 September 2001 The key objective of the research project is to uncover the principles that underlie the specification and verification of intelligent cooperative information systems (CISs). This project studies techniques for coordination, agent communication languages, commitments among agents, and commitment-based protocols. It develops both the basic concepts and their formalizations.

User Interface Softbot

Robert St. Amant

\$497,384 by the National Science Foundation

August 2000 - August 2003

This research focuses on the concept of Ibots, interface agents that interact with software applications through the graphical user interface, in essentially the same way that human users do. The proposal will produce an environment for agent exploration and evaluation, and a model of tool use in the interface. The project will produce tools that facilitate research in AI planning, user interface agents, and cognitive modeling. It will also give results of general interest to the agents community.

MIPER: Mixed-Initiative Plan Evaluation and Repair

Robert St. Amant

\$234,869 Subcontract to the University of Massachusetts, for a contract to

DARPA/Rome Laboratory,

Planning and Decision Aids.

August 1997 - May 2001

This project examines mixed-initiative planning issues in an interactive simulation of abstract physical agents. Agents generally behave in a goal-directed manner: they plan their movements, exert force on one another, cooperate in their interactions. Various scenarios establish overall goals for teams of agents, such as the destruction or

containment of opposing teams, the occupation of territory, or the capture of specific targets. The problem is integrating the user's activities with those of automated planning tools, treating the whole as a joint cognitive system. Our approach relies on novel AI planning techniques and their application through a physical schema theory.

Computational Methods in Markov Chains

William Stewart, Carl Meyer

\$333,574 by the National Science Foundation

June 1998 - May 2001

The theory and application of Markov chain techniques constitute a unifying theme in the application of mathematics to many problems in engineering, economics, physical science, and social science, and the numerical computation of stationary probabilities associated with large scale applications is a fundamental concern. This is particularly true in applications such as the modeling of flexible manufacturing systems, telecommunication systems, computer performance evaluation, biological systems, economic modeling and forecasting, and more generally, in applications where discrete models are used to understand and analyze the dynamics of large systems comprised of a collection of loosely coupled subsystems. This project emphasizes both the practical and theoretical aspects surrounding the application and implementation of algorithms for computing stationary probabilities.

Analysis and Testing of Concurrent Object-Oriented Software

K.C. Tai

\$126,590 by the National Science Foundation

September 1999 - August 2001

Due to the popularity of object-oriented technology and network applications, the use of concurrent object-oriented programs is increasing significantly.

The goal of the project is to develop efficient and effective techniques for analyzing and testing concurrent object-oriented programs. This project will investigate the following problems: (a) How to efficiently and effectively analyze and test a class that defines data to be shared by multiple processes? (b) How to efficiently and effectively analyze and test a set of classes that define processes communicating with each other? (c) How to perform incremental analysis and testing of a concurrent object-oriented program according to the relationships between classes in the program?

Girls on Track

Mladen Vouk

\$188,692 by the National Science Foundation

September 1998 - August 2001

The purpose of the Girls on Track (GoT) project is to increase Middle Grade Girls' interest in math-related careers by engaging them in computer-based mathematical

explorations of urban problems in their communities. It is a joint project between North Carolina State University (Center for Research in Mathematics and Science Education and Department of Computer Science), Meredith College, Wake County Public Schools, and North Carolina Department of Public Instruction.

Incorporating the Collaborative Software ProcessSM

Laurie A. Williams

\$108,107 by Eyecast, Inc.

September 2000 - August 2001

eXtreme Programmers (XP) have anecdotally publicized the use of pair-programming for the development of high-quality software products. Dr. Williams quantitatively supported XP's anecdotal claims of the benefits of pair-programming through an empirical study of advanced undergraduates at the University of Utah. She has also developed the Collaborative Software Process (CSP), which places the use of pair-programming in a disciplined software process. This grant proposes that a software development group in Eyecast, Inc. is educated in CSP. Process data will be collected by the software engineers in this group. This data will be used to analyze the effectiveness of pair-programming. It also proposes the development of two software tools (1) one to be used for on-demand computer-based training in the CSP and (2) a process data collection and analysis tool.

Integrating Collaborative Programming into a Disciplined Software Process: A CACC Enhancement Project

Laurie A. Williams

\$65,000 by Nortel, Inc.

November 2000 - November 15, 2001

eXtreme Programmers (XP) have anecdotally publicized the use of pair-programming for the development of high-quality software products. Dr. Williams quantitatively supported XP's anecdotal claims of the benefits of pair-programming through an empirical study of advanced undergraduates at the University of Utah. However, more evidence is necessary before large corporations, such as Nortel, are convinced of the benefits of pair-programming and are ready to embark on large-scale deployment of the technique. This grant proposes that pair programming be integrated into Nortel's development process. Process data will be collected by the software engineers practice pair programming. This data will be used to analyze the effectiveness of pair-programming. Additionally, some engineers will work in groups of three, practicing triad-programming. Their process data will also be analyzed to determine the effectiveness of this working arrangement.

Pair-Learning in Undergraduate Computer Science Education

Laurie A. Williams

\$227,000 by the National Science Foundation

January 2001 - December, 2003

Initial experimental results indicate that pair-learning, whereby two students work together at one computer, improves the success and morale of the students. We also believe that women and minority students would especially benefit from this pedagogical approach. Initial experimentation with pair-learning also reveals benefits to computer science educators. Students working in pairs are able to answer each others questions. They no longer look to the teaching staff as their sole source of technical advice; educators are no longer burdened by an onslaught of questions. Grading can be significantly reduced when two students submit one assignment. The number of cheating cases is reduced because collaboration is legitimized. The classes are calmer; the students are more satisfied and self-sufficient. These observations will be investigated further via a multiyear, extensive NSF-funded study of beginning computer science students.

CAREER: Plan-Based Integration of Control and Coherence in Intelligent Exploratory Environments

R. Michael Young

\$465,695 by the National Science Foundation

March 2001 - April 2006

Exploratory environments --- virtual worlds in which users interact with a simulation of a physical or abstract space to form, pursue and achieve their goals --- have shown great success in applications ranging from education and training to social interaction. A principal limitation of these systems, however, is that users' activities within them are typically greatly over- or under-constrained. In this project, I will develop new plan-based models for the structure of user interactions within exploratory environments. Planning techniques will be used to create novel activity within the environment that encompasses both the system-controlled characters, the environment and the actions of the user. One central hypothesis of this work is that a user's understanding of the plan-based structure of activity in a world provides scaffolding for her own exploratory efforts. Presenting the user with an environment in which change can be readily understood in terms of plans encourages the user to acquire and employ plan-based knowledge of the virtual world, to form plans that integrate with those unfolding around her and to engage in the world by executing those plans. This activity can potentially lead to more effective and engaging interaction and thus to an increased understanding of the world that the environment models.
