Computer Science Department

Professor Kay G. Schulze
Chair

During the 2005-2006 academic year, the Computer Science department faculty did significant research in the areas of software engineering, robotics, computer algebra, graphics, semantic web, and database. Two midshipmen independent research projects were conducted. Midshipmen also participated in summer internships at the Naval Research Laboratory, the National Security Agency, the Naval Surface Warfare Center, SHAPE Headquarters, Wyle Labs and the Naval Warfare Development Command. The Computer Science faculty continues to attract grant money and research opportunities from government sources.

Sponsored Research

Practical Computing with Semi-Algebraic Sets via Cylindrical Algebraic Computation

Researcher: Associate Professor Christopher W. Brown
Sponsor: National Science Foundation (NSF)

The purpose of this research was to develop algorithms and produce programs that are able to perform practical symbolic computations with semi-algebraic sets. The primary tool for this was Cylindrical Algebraic Decomposition (CAD), which provides a data-structure for explicitly representing semi-algebraic sets. CAD is a powerful and general tool, which has been implemented several times and already shown itself to be of more than just theoretical value. However, the consensus in the literature is that interesting application problems cannot be solved by current CAD programs within a reasonable amount of time and space. The three principal goals of the project were to extend the theory of CADs to deal efficiently with the types of semi-algebraic sets that arise in applications, to produce efficient algorithms and programs that compute with CADs, and to promote the use of CADs to solve problems in other disciplines.

Imitation Learning of Goal Sequences in Mobile Robots

Researcher: Associate Professor Frederick L. Crabbe, IV
Sponsor: Office of Naval Research (ONR)

Most studies of imitation learning in robots or other artificial agents concentrate on enabling a robot to learn to repeat individual low-level actions made by a teaching agent. This project continued to expand imitation learning to enable a student robot to learn to identify and perform sequences of high-level sub-goals, each made up of multiple low-level actions. For example, a student agent would follow a teacher agent as the teacher performed a task in order to learn to perform that same task. When the teacher completed each sub-goal in the sequence, the student learned the purpose of each sub-goal as well as when to perform it in the sequence. Work focused on enabling a learning robot to segment the stream of motor commands performed by the teaching robot into conglomerate actions. The learning robot used Bayesian techniques to determine the relative probabilities of various motions in order to determine if one action had ended and another had begun.

Dealing with Uncertainty when Agents Meet Real People

Researcher: Assistant Professor Luke K. McDowell
Sponsor: Naval Academy Research Council (NARC)

To be effective, semantic agents must be able to easily and effectively interact with a wide range of participants in order to achieve some goals. Unfortunately, such agents face many problems, including the need to interact with human participants that do not have a suitable agent, the tendency of participants to respond never or

more than once, and the possibility that unexpected responses may force a mid-course change to the agent's goals. In this work, we used our deployed system for semantic email agents (E-Agents) to address these general challenges. In particular, we demonstrated how messages that are actionable by either humans or agents can deal with the agent-less participant problem. We then introduced a fundamental notion of "eager responses" and explained how such responses could address the problems of time-varying responses and goals. We demonstrated that reasoning about such responses is intractable in the worst case, but can often be done in polynomial time. Finally, we discussed additional semantic reasoning procedures to predict and interpret participant responses. These results both greatly increased the usefulness of E-Agents as well as identified a number of usability and functionality issues that apply to a much broader range of semantic agents.

High Energy Laser Weapons: Modeling and Simulation Phase II, Modeling Development.

Researcher: Associate Professor Donald M. Needham
Sponsor: The High-Energy Laser Joint Technology Office

Development of one or more end-to-end simulations of laser system performance and effectiveness is fundamental to the goals of the High Energy Laser Joint Technology Office's (JTO) modeling and simulation (M&S) program. One such simulation prospective follows the physics and engineering of a ship's energy starting with energy conversion from the ship's fuel through electrical power generation, electrical generation of laser light, beam transport, etc, and ending with target damage assessment. A fundamental step, proposed by this research, in the process of developing a domain-reusable, end-to-end high energy laser simulation, was to examine an existing model, and focus on the identification of a software system architecture designed to support the inclusion of modeling components needed for the various portions of an end-to-end simulation. The steps in this process included acquiring a suitable high energy laser model, and developing an architecture that isolated and minimized couplings between the software components that implemented the model. Analysis of the resulting architecture focused on increasing its potential for domain-specific reuse so that suitable models for the remaining portions of the end-to-end simulation could be smoothly integrated. Additionally, this work provided guidance to high energy laser modelers on developing software components that implement their models so as to ease integration of their models within an end-to-end simulation.

A Tactical Data Simulator for Network Centric Training Operations at the Naval Academy

Researcher: Assistant Professor David J. Stahl, Jr.
Sponsor: Office of Force Transformation (OFT)

The Yard Patrol (YP) Tactical Data Simulator (YP-TDS) is an inexpensive commercial-off-the-shelf hardware and software system for simulating sensors, platforms and weapons in real-time across a wide-area network using HF packet radio. The system integrates real sensors with simulated weapons, sensors and ships by means of a multi-terminal command-and-control tactical display having appearance and functionality similar to that of existing warship systems – but intended for use aboard non-commissioned, non-warship vessels in an educational environment. YP-TDS represents an initial effort at embodying the concept of Network Centric Warfare not just in theory, but in actual hands-on practice at USNA. A multiple-actor battle simulation framework was implemented and initial operational testing was successfully completed. The system was extended to support use in the classroom. Additional software development and creation of pedagogical tools and materials is in progress.

Independent Research

An Investigation into Compromise Actions in Behavior-Based Robotic Systems

Researcher: Associate Professor Frederick L. Crabbe, IV

This project examined the nature of compromise actions in behavior-based systems. Using Value-Iteration to approximate solutions to the Bellman equation, it determined the comparative advantage to selecting compromise actions when pursuing two appetitive goals. It then used Information Gain to determine in what scenarios compromise behavior provided the greatest benefit.

A Comparative Performance Study of P2P Index Structures for Range Queries

Researcher: Assistant Professor Adina N. Crainiceanu

Peer-to-peer (P2P) systems provide a robust, scalable and decentralized way to share and publish data. The purpose of this project was to implement and compare the performance of different index structures. We implemented P-Ring, Skip Graphs, Online Balancing, and Chord and performed preliminary experiments to study their performance. We plan on expanding the study to include Mercury, a representative P2P range index structure, and to evaluate the performance of the implemented index structures in a large area distributed network.

Automated, Domain-focused Information Extraction from the Web

Researcher: Assistant Professor Luke K. McDowell

The Semantic Web's need for machine understandable content has led researchers to attempt to automatically acquire such content from a number of sources, including the web. To date, such research has focused on document-driven systems that individually process a small to moderate size set of documents, annotating each with respect to a given ontology. This work introduced *OntoSyphon*, an alternative that strives to more fully leverage existing ontological content while scaling to extract comparatively shallow content from billions of documents. *OntoSyphon* operates in an ontology-driven manner: taking any ontology as input (with or without instance data). *OntoSyphon* uses the ontology to specify web searches that identify possible semantic instances, relations, and taxonomic information. Information from the ontology, together with redundancy in the web, was then used to automatically verify these candidate instances and relations, enabling *OntoSyphon* to operate in a fully automated, unsupervised manner. A prototype of *OntoSyphon* was fully implemented and we presented experimental results that demonstrate substantial instance learning in a variety of domains based on independently constructed ontologies. We also introduced new techniques for improving instance verification, both with and without ontological background knowledge, and demonstrated empirically that they improved upon the best previously known techniques.

Exploring Dimensionality Reduction for Text Mining and Literature-Based Discovery

Researcher: Assistant Professor Luke K. McDowell

The world is awash in data. This information overload is made all the more difficult by the great variety of information that must be considered, and by the need to perform analysis very quickly. The free form nature of this information necessitates the use of text mining methodologies for analysis. Text mining is the process of automatically extracting information from unstructured text. The primary challenge with text mining is the complexity of a document and the process by which this complexity is simplified into a more manageable representation. A number of techniques offer great promise for effectively reducing vast amounts of information while maintaining the most interesting details, but have yet to be adequately evaluated for text mining. This work proposed to evaluate the effectiveness of several linear and nonlinear methods for dimensionality reduction as they relate to literature-based discovery and other text mining processes.

Project Angel Fire

Researchers: Captain Paul Young, USN, Assistant Professor David J. Stahl, Jr.,
and Captain Brian Balazs, USAF

Project Angel Fire, sponsored by Los Alamos, is a rapid technology development project designed to provide geo-located situational awareness to the war fighter ("urban radar"). Angel Fire is envisioned to provide distribution of very large area sensor imagery from potentially many types of sensors using "ordinary" computer hardware to the soldier on the ground in addition to analysts and command center personnel. Angel Fire is expected to use available links and storage at acceptable levels of service, including providing distribution over disadvantaged networks.

Bowman Scholar Project

Applying Semantic Web Concepts to Support Net-Centric Warfare and the GIG Using TAML

Researcher: Midshipman 1/C Candace M. Childers, USN
Adviser: Captain Paul Young, USN
Sponsor: Bowman Scholar Program (BSP)

The ability to analyze data quickly and transform it into actionable information is vital for information superiority. However, the amount of available data is increasing and the time to make decisions is decreasing. There is too much data for humans to sift through and filter for decision making, so computer automation is necessary. The Extensible Markup Language (XML) offers a partial solution by providing a syntactic standard for data exchange. The Tactical Assessment Markup Language (TAML) is an XML vocabulary for exchanging undersea warfare tactical data. However, the meaning or semantics of the data is unknown to the machine processing the data.

The Semantic Web is a set of technologies designed to add semantic information to data for machine processing. The technologies consist of several components, including a common syntax for data exchange, common semantic representation, and a common ontology language. Reasoning engines apply algorithms to the data to infer useful information and present it to decision makers. Sophisticated Semantic Web tools and techniques are rapidly emerging. This research investigated adding stronger semantic content through application of Semantic Web technologies to XML-based languages such as TAML. The lessons learned will help enable systems to extract useful, actionable information from a number of distributed, autonomous, heterogeneous information sources and bring the armed forces closer to a knowledge-aware Global Information Grid (GIG).

Midshipman Research Course Projects

A Blue Force Tracking System for Midshipmen (BFT-M)

Researcher: Midshipman 1/C Paul K. Evans, USN
Adviser: Assistant Professor David J. Stahl, Jr.

This project explored the feasibility of networking Windows CE based handheld devices using inexpensive off-the-shelf hardware and software systems to provide midshipmen with a tactical training system simulating the FBCB2 system - Force XXI Battle Command, Brigade and Below ("Blue Force Tracking"). In conjunction with the YP Tactical Data Simulator, “Midshipman Blue Force Tracking” is intended to be used as a pedagogical tool for educating midshipmen in the concepts of Network Centric Warfare and operations. Published as United States Naval Academy Computer Science Department Technical Report, USNA-CS-TR-2005-08, 2005.

IPv6 Testing

Researcher: Midshipman 2/c Christopher B. Landis, USN
Advisers: NSA Visiting Professor Thomas R. Hendricks and Mr. Kevin Keeley (NSA)

The Department of Defense (DoD) is committed to transitioning to Internet Protocol version 6 (IPv6) by 2008. Transitional working groups are under way to develop plans to implement IPv6 on Government networks. The DoD has formed working relationships to conduct testing and share information. This project performed some of the testing. Published as United States Naval Academy Computer Science Department Technical Report, USNA-CS-TR-2006-02, 2006.

Professional Publications

Journal (Refereed) Manuscripts

BROWN, Christopher W., Associate Professor, El Kahoui, M'hammed, Novotni, Dominik, and Weber, Andreas, "Algorithmic Methods for Investigating Equilibria in Epidemic Modeling," *Journal of Symbolic Computation*, 2006, forthcoming, in press.

> The calculation of threshold conditions for models of infectious diseases is of central importance for developing vaccination policies. These models are often coupled systems of ordinary differential equations, in which case the computation of threshold conditions can be reduced to the question of stability of the disease-free equilibrium. This paper showed how computing threshold conditions for such models can be done fully algorithmically using quantifier elimination for real closed fields and related simplification methods for quantifier-free formulas. Using efficient quantifier elimination techniques for special cases that have been developed by Weispfenning and others, we also computed whether there are ranges of parameters for which sub-threshold endemic equilibria exist.

CRABBE, Frederick, L., IV, Associate Professor, "Unifying Undergraduate Artificial Intelligence Robotics: Layers Of Abstraction Over Two Channels," *AI Magazine*, Vol. 27, No. 1, 2006, pp 23-37.

> From a Computer Science and Artificial Intelligence perspective, Robotics often appears as a collection of disjoint, sometimes antagonistic sub-fields. The lack of a coherent and unified presentation of the field negatively impacts teaching, especially to undergraduates. The paper presented an alternative synthesis of the various sub-fields of Artificial Intelligence robotics, and showed how these traditional sub-fields fit in to the whole. Finally, it presented a curriculum based on these ideas.

Van Lehn, K., Lynch, C., SCHULZE, Kay G., Professor, Shapiro, J. A., SHELBY, Robert H., Professor Emeritus, Taylor, L., TREACY, Donald J., Professor, Weinstein, A., and WINTERSGILL, Mary C., Professor, "The Andes Physics Tutoring System: Lessons Learned," *International Journal of Artificial Intelligence and Education*, Vol. 15, No. 3, 2005, pp.1-47.

> The Andes system demonstrates that student learning can be significantly increased by upgrading only their homework problem-solving support. Although Andes is called an intelligent tutoring system, it actually replaces only the students' pencil and paper as they do problem-solving homework. Students do the same problems as before, study the same textbook, and attend the same lectures, labs and recitations. Five years of experimentation at the United States Naval Academy indicate that Andes significantly improved student learning. Andes' key feature appeared to be the grain-size of interaction. Whereas most tutoring systems have students enter only the answer to a problem, Andes has students enter a whole derivation, which may consist of many steps, such as drawing vectors, drawing coordinate systems, defining variables and writing equations. Andes gives feedback after each step. When the student asks for help in the middle of problem-solving, Andes gives hints on what's wrong with an incorrect step or on what kind of step to do next. Thus, the grain size of Andes' interaction is a single *step* in solving the problem, whereas the grain size of a typical tutoring system's interaction is the *answer* to the problem. This report was a comprehensive description of Andes. It described Andes' pedagogical principles and features, the system design and implementation, the evaluations of pedagogical effectiveness, and our plans for dissemination.

Conference Proceedings

Anderson, Richard, Anderson, Ruth, and MCDOWELL, Luke K., Assistant Professor, and Simon, Beth, "Use of Classroom Presenter in Engineering Courses," *Proceedings of the 35th ASEE/IEEE Frontiers in Education Conference*, 2005, pp. T2G-13 – T2G-18.

> Instructors of computer science, engineering, and technology courses often teach using prepared slides displayed with a computer and data projector. While this approach has advantages, disadvantages can be severe -- including limited flexibility in delivery of material, which hinders instructors' spontaneous adaptation to students. This paper described *Classroom Presenter*, a Tablet PC-based presentation system that (1) allows instructors to regain flexibility in presentation using Tablet PC-supported digital inking in an application specifically designed for lecturing use, (2) supports role-specific views and interaction mechanisms in class, and (3) introduces novel affordances including the use of wireless technology and student devices to support active learning in the classroom. As of spring 2005, *Classroom Presenter* has been used in over 200 courses at more than 13 institutions. Here we took advantage of our large user-base to discuss *Classroom Presenter*'s impact on student learning and summarized best practices for lecturing with this technology.

BROWN, Christopher W., Associate Professor, and McCallum, Scott , "On Using Bi-equational Constraints in CAD Construction", *Proceedings of the 2005 International Symposium on Symbolic and Algebraic Computation (ISSAC '05*), 2005, pp. 76-83.

> This paper introduced an improved method for constructing cylindrical algebraic decompositions (CADs) for formulas with two polynomial equations as implied constraints. The fundamental idea is that neither of the varieties of the two polynomials is actually represented by the CAD the method produces, only the variety defined by their common zeros is represented. This allowed for a substantially smaller projection factor set, and for a CAD with many fewer cells.

CRABBE, Frederick L., IV, Associate Professor, "On Compromise Strategies for Action Selection with Proscriptive Goals," *Proceedings of Modeling Natural Action Selection*, 2005, pp 24-31.

> Among many properties suggested for action selection mechanisms, one prominent one is the ability to select compromise actions, i.e. actions that are not the best to satisfy any active goal in isolation, but rather compromise between the multiple goals. This paper performed an analysis of compromise actions in situations where the agent had one proscriptive goal. It concluded that optimal compromise behavior looks quite different from what was expected, and, while optimal compromise actions are beneficial to an agent, the benefit is often small compared to greedy algorithms. It went on to suggest that much of the discussion about compromise behavior is the result of an equivocation on its definition, and it proposed a new compromise behavior hypothesis.

DELOOZE, Lori, Assistant Professor, "He Who Does Not Learn History Is Doomed to Repeat It", *Proceedings of the 6th Conference on Information Technology Education*, 2005, pp. 329–332.

> Although Information Technology (IT) seems to be changing at a shocking rate, common themes can be observed. These themes have been repeated throughout history. IT students should have a solid foundation of these historic trends and ongoing evolutionary progress before beginning the bulk of their studies in the field of computing. The course, IT Revolutions Past and Present, at the United States Naval Academy is required of IT students and sets the context for their coursework. Students explored the impact of both successes and failures in the field and analyzed the value that IT professionals bring to commercial and government organizations.

DELOOZE, Lori, Assistant Professor, "Minimal UML Diagrams for a Data-driven Web Site", *Proceedings of the 6th Conference on Information Technology Education,* 2005, pp. 229-232.

> The Unified Modeling Language (UML) was developed for visually expressing the models of software-intensive systems. The extensive UML modeling process is overwhelming for undergraduate students who are creating rather small data-driven web sites. We identified a minimal subset of UML diagrams that could be used to effectively model small, well-defined projects that used a MySQL database with PHP scripting. Sixty-five students successfully completed an undergraduate Information Technology Capstone course using this proposed methodology while developing one or two incremental prototypes and a final, self-documented working product. The step-by-step process encouraged students to follow programming "best practices", including performing planning, analysis, and design before attempting to write code. Instructors also required weekly status reports and updated project plans to make the experience as realistic as possible. The general development methodology was found to be extremely successful.

DELOOZE, Lori, Assistant Professor, "Creating a Microcosm of Industry in the Classroom," *Proceedings of the 6th Conference on Information Technology Education,* 2005, pp. 217-221.

> Electronic Commerce has challenged traditional business thinking. Electronic Businesses are expected to be always available, from anywhere in the world, rich in content and experience, customizable to the user, and conform to technology standards. Features available in such an enterprise must support both Business-to-Customer (B2C) and Business-to-Business (B2B) electronic commerce. We have designed an innovative course that teaches students how to develop an organization's Enterprise Architecture that becomes part of a microcosm of the consumer entertainment industry in the classroom. Seven groups represent each of three levels of businesses within the industry (manufacturing, distribution, and retail) and together they develop an online community to sell movies, books and music to the general public - a la Amazon.com. Students were given three major projects designed to lead to a fully-functioning, integrated industry architecture: an enterprise portal, internal messaging system and B2B communication via XML. Because students were overwhelmed by the amount of work required to complete the projects, we plan to break the tasks into weekly lab assignments for the next offering that build up to the overall working architecture.

MCDOWELL, Luke K., Assistant Professor, "Preparing Semantic Agents for an Unsuspecting World," *First International Symposium on Agents and the Semantic Web (AAAI Fall Symposium Series)*, 2005, pp. 103-105.

> To be effective, semantic agents must be able to easily and effectively interact with a wide range of participants in order to achieve some goals. Unfortunately, such agents face many problems, including the need to interact with human participants that do not have a suitable agent, the tendency of participants to respond never or more than once, and the possibility that unexpected responses may force a mid-course change to the agent's goals. In this paper, we used our deployed system for semantic email agents (E-Agents) to address these general challenges. In particular, we demonstrated how messages that are actionable by either humans or agents can deal with the agent-less participant problem. We then introduced a fundamental notion of "eager responses" and explained how such responses can address the problems of time-varying responses and goals. We demonstrated that reasoning about such responses was intractable in the worst case, but could often be done in polynomial time. Finally, we discussed additional semantic reasoning procedures to predict and interpret participant responses. These results both greatly increased the usefulness of E-Agents as well as identified a number of usability and functionality issues that apply to a much broader range of semantic agents.

NEEDHAM, Donald M. Associate Professor and Jones, Sean A., (Trident Scholar 2005), "Analyzing Software Fault Trees via a Key Node Metric," *Proceedings of the Software Metrics European Forum, (SMEF'06)*, 2006, pp. 237-246.

Complex software systems for business critical and safety critical applications require tools for business performance managers to use, especially in software systems in which failure leads to major economic failure or loss of life. The software engineering community stands to benefit from metrics, analysis tools, and techniques that address software system safety from a design perspective. The design-time use of software fault trees in representing the structure of a software system allows designers to focus on business and safety critical aspects of software during early development stages thereby allowing business performance managers to provide cost-effective oversight of software development.

This paper applied a technique for evaluating systems through the analysis of software fault tree "key nodes" that required multiple inputs to fail before the failure being considered would propagate to other parts of the software system. The metric provided business managers with a tool through which to control costs during software development. A heuristics-based key node metric was presented, providing a design tool with which to compare fault trees without requiring a priori knowledge of component reliability. The ability of the metric to be applied without requiring component reliability knowledge allowed the metric to be used at design time where component reliability values for a type of component are often unknown. The key node metric allowed designers to proactively improve the business or safety critical aspects of a system before final component selection or the completion of component reliability studies. The paper provided an application of the key node metric, and discussed the results of applying the metric to a product line represented by a set of software design mutations.

NEEDHAM, Donald M. Associate Professor, "Interdisciplinary Teams for Software System Development," *Proceedings of the 2005 International Conference on Frontiers in Education: Computer Science and Computer Engineering (FECS'05)*, 2005. pp. 10-16.

Embedded software systems typically require the combined efforts of designers and engineers with a wide variety of software, hardware, and system interfacing backgrounds. A crucial part of embedded software development centers on actively involving interdisciplinary teams in all stages of the software development process. This paper presented the design and development activities of several projects undertaken by interdisciplinary teams of undergraduate students. Each team included students from a variety of academic backgrounds, including aerospace engineering, computer science, mechanical engineering, oceanography, and systems engineering. The teams developed, tested, and deployed embedded control systems for an autonomous underwater vehicle, and several unmanned aerial vehicles. This paper discussed lessons learned from working with interdisciplinary teams of students on embedded systems, the influence of team organization on the development process, and the effect of software engineering-oriented decisions such as development methodology, prototyping, and testing techniques.

NEEDHAM, Donald M. Associate Professor, and SCHULZE, Kay G. Professor, "Discovering Student Outcomes," *Proceedings of the 2005 International Conference on Frontiers in Education: Computer Science and Computer Engineering (FECS'05)*, 2005. pp. 94-99.

The Computing Accreditation Commission of ABET, Inc. requires computer science programs to identify program outcomes that are uniquely determined by the programs themselves. This paper discussed the process used by an accredited computer science department to establish, measure, and analyze program outcomes. Program outcomes were considered from the perspectives of both students and faculty, and a tool that can assist with program outcomes measurement was examined.

Linga, Prakash, CRAINICEANU, Adina N., Assistant Professor, Gehrke, Johannes, Shanmugasundaram, Jayavel, "Guaranteeing Correctness and Availability in P2P Range Indices," *Proceedings of International Conference on Management of Data (SIGMOD 2005)*, 2005, pp. 323-334.

> New and emerging P2P applications require sophisticated range query capability and also have strict requirements on query correctness, system availability and item availability. While there has been recent work on developing new P2P range indices, none of these indices guarantee correctness and availability. In this paper, we developed new techniques that can provably guarantee the correctness and availability of P2P range indices. We developed our techniques in the context of a general P2P indexing framework that could be instantiated with most P2P index structures from the literature. As a specific instantiation, we implemented P-Ring, an existing P2P range index, and showed how it could be extended to guarantee correctness and availability. We quantitatively evaluated our techniques using a real distributed implementation.

Vanlehn, K., Lynch, C., SCHULZE, Kay G., Professor, Shapiro, J. A., SHELBY, Robert H., Professor Emeritus, Taylor, L., TREACY, Donald J., Professor, Weinstein, A., and WINTERSGILL, Mary C., Professor, "The Andes physics tutoring system: Five years of evaluations," in G. I. McCalla and C.-K. Looi (Eds.), *Proceedings of the Artificial Intelligence in Education Conference*, 2005, pp. 678-685.

> Andes is a mature intelligent tutoring system that has helped hundreds of students improve their learning of university physics. It replaces pencil and paper problem solving homework. Students continue to attend the same lectures, labs and recitations. Five years of experimentation at the United States Naval Academy indicate that it significantly improved student learning. This report describes the evaluations and what was learned from them.

Technical Reports

CRAINICEANU, Adina N., Assistant Professor, "Answering Complex Queries in Peer-to-Peer Systems", Doctoral Dissertation, Cornell University, May 2006, pp. 1-171.

> Peer-to-peer (P2P) systems provide a robust, scalable and decentralized way to share and publish data. However, most existing P2P systems only support equality or keyword search queries. We believe that future P2P applications, such as digital libraries, resource discovery on a grid, or military applications, will require more complex query functionality, as users will publish semantically rich data.
>
> Towards the goal of supporting complex queries in P2P systems, this dissertation focused on developing index structures that allowed fast access to distributed data. We introduced first a general modular indexing framework that identified and separated the different functional components of a P2P index structure. One of the benefits of such a framework was that it allowed reusing the existing algorithms for different components rather than implementing everything anew. We were thus able to concentrate on creating algorithms that provided new functionality.
>
> We introduced P-Tree, a distributed fault-tolerant index structure. P-Trees support range queries in addition to equality queries in P2P systems in which each user (peer) publishes a small number of data items. We then introduced P-Ring, a novel index structure based on our framework, that supports both equality and range queries, is fault-tolerant, efficiently supports large sets of data items per peer (as opposed to P-Tree), and provides guaranteed logarithmic search performance in a stable system. We evaluated the performance of P-Ring and compared P-Ring with three other P2P index structures, Chord, Skip Graphs, and Online Balancing, implemented in the context of our indexing framework. Our performance results indicated that P- Ring outperformed Skip Graphs in terms of both query and maintenance cost. P-Ring offered a better load balance than Online Balancing, at a lower cost. P-Ring outperformed Chord in terms of search performance and supported a larger class of queries, with a low overhead.

NEEDHAM, Donald N., Associate Professor and Jones, Sean A., (Trident Scholar 2005), "Software Fault Tree Key Node Metric Test Cases," United States Naval Academy Computer Science Department Technical Report, USNA-CS-TR-2006-01, 2006.

This report contained 70 sets of software fault trees used to test a software fault tree key node safety metric.

Presentations at Professional Meetings and Conferences

CRABBE, Frederick L. IV, Associate Professor, "On Compromise Strategies for Action Selection with Proscriptive Goals," Modeling Natural Action Selection, Edinburgh, Scotland, July 2005.

CRABBE, Frederick L. IV, Associate Professor, "Compromise Strategies for Action Selection," Intelligent Systems Program, University of Pittsburgh, Pittsburgh, PA, March 2006.

DELOOZE, Lori, Assistant Professor, "Applying Soft Computing Techniques to Intrusion Detection," Cyber Security & Information Infrastructure Research Workshop, Oak Ridge National Laboratory, Joint Institute for Computational Sciences, Oak Ridge, TN, May 2006.

MCDOWELL, Luke K., Assistant Professor, "Preparing Semantic Agents for an Unsuspecting World," First International Symposium on Agents and the Semantic Web (AAAI Fall Symposium Series), November 2005.

MCDOWELL, Luke K., Assistant Professor, and LOGUE, Thomas, Commander, USN, "Using Tablet PCs to Enhance Classroom Learning,", Computer Science Department Seminar, Annapolis, MD, November 2005.

NEEDHAM, Donald, M., Associate Professor, "Interdisciplinary Teams for Software System Development," International Conference on Frontiers in Education: Computer Science and Computer Engineering (FECS'05), Las Vegas, NV, June 2005.

NEEDHAM, Donald, M., Associate Professor, "Discovering Student Outcomes," International Conference on Frontiers in Education: Computer Science and Computer Engineering (FECS'05), Las Vegas, NV, June 2005.

NEEDHAM, Donald, M., Associate Professor, "Analyzing Software Fault Trees via a Key Node Metric", Software Metrics European Forum, (SMEF'06), Rome, Italy, May 2006.