

Welcome to Black History Month! February 2008

Dear Students:

Our once-per-year opportunity to celebrate Black History is upon us and we need to make the most of it. For some, history is distant and abstract, consisting of hard to believe events and trends often analyzed in dry prose. Survival in the present may mean dissociation with the past; however, perspective on the past is important to survival in the future.

Since a culture is more than just a collection of individuals, we must strive to preserve those aspects that manifest its core vitality. Preservation of culture requires the retention of customs and beliefs critical to the identity of the whole. As members of a minority culture within the USA how do we develop personal identities that enhance the collective image of Black America, at a time when media stereotypes paint a narrow and negative portrait.

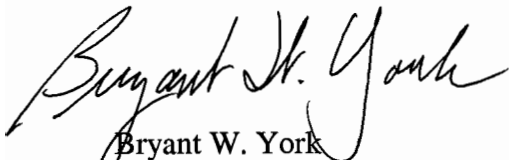
In the pages that follow I present some of the accomplishments of 10 young African American men who are doing advanced research in computer science. A separate document hailing the achievements of young African American women researchers in computer science is being developed by a colleague.

Let's start by saying what these young men are not! They are not gang bangers; they are not pimps; they are not drug dealers or drug addicts; they are not even professional athletes or entertainers. They are highly accomplished, young, African American Ph.D. researchers in computer science.

These brief biographies and research statements just begin to tell the amazing stories behind their rise to technological prominence.

For Black History month 2008 let's acknowledge and celebrate the contributions to American Society and to the core vitality of African American culture of these young men.

Sincerely



Bryant W. York
February 2, 2008

Dr. Brian Blake, Georgetown University

<http://www.cs.georgetown.edu/~blakeb/website/index.htm>



M. Brian Blake is a native of Savannah, Georgia graduating from Benedictine Military High School. He graduated from Georgia Institute of Technology with a Bachelor of Electrical Engineering in 1994. His first industry position was with G.E. Aerospace/ Lockheed Martin Real3D as a Technical Leadership/Edison Engineer, which allowed him to complete a Master's degree while working as a Software Engineer completing 4 rotations. His rotations were in ASIC design, software development for embedded controllers, and C++ development. Upon graduating from the Technical Leadership Program, he took a position as Software Integration Consultant, with Lockheed Martin Integrated Business Solutions in Atlanta. While continuing the degree part-time at Mercer University with project work at Georgia Institute of Technology, he was a consultant with Lockheed Martin and held a part-time position as Instructor in the training division, Lockheed Martin Advance Concepts Center. He acted as a trainer-trainee for the Object-Oriented Analysis and Design and Design Pattern courses. He graduated with a degree from Mercer University with a Master of Science in Electrical Engineering and a minor in Software Engineering in 1997. Dr. Blake took a nontraditional path by completing his PhD while working full-time as a software engineer. He started the PhD in Information Technology and Engineering/Computer Science at George Mason University. In April 1998, he took a position as a Sr. Computer Scientist with Trident Data Systems (now General Dynamics) designing and developing Windows applications in the intelligence domain. He graduated with his doctorate in Information Technology/Computer Science from George Mason University in the fall of 2000.

Currently he is an Associate Professor and Chair in the Department of Computer Science at Georgetown University where he was one of the youngest African-Americans in history to get tenure in computer science department. Dr. Blake conducts applied research in the development of automated approaches for the sharing of information and capabilities across organizational boundaries, sometimes referred to as enterprise integration. With respect to this area of interest, his investigations cover the spectrum of software engineering: design, specification, proof of correctness, implementation/experimentation, performance evaluation, and application. He has published over 80 journal articles and refereed conference papers in the areas of service-oriented computing, intelligent agents and workflow, enterprise systems integration, component-based software engineering, distributed data management, and software engineering education. He has been the recipient of two best paper awards. His work has been funded over \$5 million from the National Science Foundation, DARPA, Federal Aviation Administration, the MITRE Corporation, Air Force Research Lab, SAIC, and the National Institute of Health.

Dr. Blake was recognized as the "Most Promising Scientist" in 2003 by the Career Communications group, and in 2007 as one of 10 Emerging Scholars in *Diverse Magazine*. In 2003, he was recognized as an Outstanding Alumnus from the Department of Software

Engineering at George Mason, just three years after graduation. Dr. Blake his highly active in the recruitment and retention of underrepresented minorities into science and technology careers.

Research Statement

Although the notion of computer software have been around for over 70 years, the methodology, design, and implementation (i.e. software engineering) of software applications still rests in the hands of humans. As a practicing software engineer for 7 years, it occurred to me that the best architected systems were those developed by the most *experienced* software engineers. Couldn't that experience be captured and used to design and develop *software that can write other software*? As a result, I entered graduate school with the initial motivation of implementing artificially intelligent software that can detect patterns in existing software and use that knowledge to build new software. It turns out that my initial intentions were quite ambitious considering the variability of programming languages (much similar to variability in the human language). However, with the inception of modular software components (more recently known as software *services*) in the mid-90s, a more reasonable goal was to develop approaches where higher-level software components could be combined and *composed* to create new applications. My specific emphasis has been in connecting business services across organizations.

My research career has been a quest to develop an intelligent software agent or group of intelligent agents (i.e. software robots) that can be used to clone my own software engineering experience in developing distributed software systems. In order for this vision to be realized, there are several milestones, that must be accomplished, many of which have been addressed within my career research projects.

- **Humans must be able to transfer knowledge about software development to intelligent software agents.** In order to accomplish this milestone, one must ask the question how are software engineering instructions being passed *among humans* much less software agents? It has been commonly accepted that visual models are the most effective way to capture software engineering artifacts. My initial work was toward the architecture of an intelligent software agent that could receive and process workflow-based visual specifications and automatically compose software systems. A unique innovation within my work has been the design, experimentation, and evaluation of the use of Unified Modeling Language (UML)-based models to capture requirements for the development of new systems as interpreted by software agents. My latest work in this area introduces software development lifecycles where human software engineers and intelligent software agents work together. Although I focus on the exploitation of agents, the new software lifecycles generated in my work are relevant to any set of actors whether they are a group of agents or a group of human software developers. We have also produced simulation software that provides decision support when evaluating the system design using the visual models.
- **Intelligent software agents must be able to autonomously discover information about available software capabilities while adapting to the on-going changes in the environment.** Given an environment where software is advertised and freely available over the Internet, software agents must be able to find existing software as a prerequisite to creating new applications. The standard software interfaces promoted in service-oriented

computing paradigms provide the foundation, however software is fundamentally different since it is ultimately defined in different organizations. As such, messages that pass into and out of software services are heterogeneous. In some cases, the same phrase can have different meanings depending on the context of the organization that created it. Leading approaches to data and software service integration suggest the use of ontology or a semantic description of the words (even that of my own work to be described in the next bullet), but these approaches tend to be tedious and the resulting files are bulky. A unique innovation in my work is the use of syntactic, natural language-based approaches to determine when the underlying parameters of software services are equivalent. This approach does not require additional metadata. A feature of our approach is that the algorithm evolves as intelligent software agents encounter and assimilate more software. Indeed, this is a stride toward my first graduate idea of software that examines existing software to determine new software engineering plans. In fact, our work tries to interpret the intentions and tendencies of the original software developers and use that knowledge to predict what other developers might do. In related research, we have exploited the organizational context of the user to recommend services for new business capabilities.

- **Intelligent software agents must be able to connect and compose software capabilities, in real time, based on user needs and limited information.** Although many approaches to the composition of software consider simple, sequential connections, in the real world, business processes are complex and intricate. Central to all of my work has been the notion of workflow which facilitates the modeling and implementation of complex processes. Perhaps the centerpiece of my research has been my contributions to the composition of software using workflow paradigms. In my own work, I have customized workflow languages for use in service composition scenarios. I have also integrated semantic descriptions to assist composition. My work has contributed heuristics and intelligent software agent approaches to compose software. In 2004, I extracted heuristics from my research projects and created a list of service composition challenges. In 2005, I founded and organized a competition where peer researchers could showcase their service composition approaches side-by-side. This competition has gained wide notoriety and the heuristics from my lab have been the stimulus for many researchers who investigate service composition.
- **Some domains require general software engineering and data management architecture to promote data sharing in addition to automated software.** Since data is always evolving, there is a need for software that can adapt as data changes. A contribution of my work has been the development of distributed architectures that allow organizations to share their data automatically. An innovation in my work has been the integration of data extraction and management services to enable the processing of information across distributed organizations. A feature in my work is the use of workflow to integrate data management services. This work was the subject of a patent application and has been deployed to many government organizations dealing with mission-critical applications (e.g. Department of Justice, Department of Defense, and Federal Aviation Administration).
- **If my vision is realized and software agents begin to build new systems, then soon agents will also have to train human users how to operate these new systems.** The most recent thread of my research has been the use of agents and workflow to develop systems that

train users. An innovation of my work to date has been the development of an intelligent agent that can monitor the performance of a user and customize the best training approach for that human user. Initial results show that we have been able to enhance human learning performance by more than 15% using our customization of the pedagogical approach called scaffolding. Connecting this agent learning work with our service-oriented computing work, the intention is to allow agents to train humans in distributed locations while acquiring training information through service discovery. In collaboration with Georgia Institute of Technology, our agent learning approaches will be incorporated into robot control mechanisms. In the past 3 years, we have acquired \$1.5 million in awards funded by the National Science Foundation. Eighty percent of this funding will be carried out between 2007 and 2010, so we are anticipating several innovations in the next 5 years in this new area for my lab.

- **My future research plans.** In the next five years, I plan to continue to progress in the area of service-oriented computing as my main line of research. I am a co-PI of two large cooperative grants over the next 4 years with over \$7 million of total funding and \$700K coming to my research lab. One project is looking at the integration of software services as a storefront to allow novice users to work within high performance computing (HPC) environments (i.e. large environments of connected computers that allow users to compute complex tasks rapidly). We like to call this an *HPC-Kinkos*. As the service-oriented computing expert within the project, I am a part of a new consortium working on a second phase of DARPA funding towards this end. The other project, to be funded by the Air Force Office of Scientific Research, exploits both my service-oriented computing work and distributed data management work to develop an architecture that will allow government organizations of various classification levels to seamlessly share data. I will also continue my work in the use of intelligent agents to train human users. As previously mentioned, my lab has two NSF grants to conduct this research over the next 3 years. To date, my sponsored research supports 2 postdoctoral scholars, 3 graduate students, and 6 undergraduates over the next 3 years (not including the pending funding). With current funding and the potential success of pending awards, I plan to establish a Center for Services Research that encapsulates all of my work.

Dr. Jayfus Doswell, Juxtopia, Inc.

<http://www.juxtopia.com>



Jayfus T. Doswell is founder, president, and CEO of Juxtopia, LLC, a privately held biomedical and information technology (BIO-IT) company and an innovative leader in human performance monitoring products and services. Juxtopia's innovative products are designed to integrate into a human's daily routine and provide services to improve human health and learning for a lifetime. Juxtopia, along with its research alliances, conducts empirical multidisciplinary research on how convergent emerging technologies can quantitatively improve human health and learning performance.

Dr. Doswell is also founder and board chairperson of the Juxtopia Group, Inc., a 501 c (3) non-profit organization established in 2000. The Juxtopia Group develops and evaluates learning technology tools for distribution into informal learning environments for the purpose of achieving the following goals:

- Reduce the critical learning gap of underserved and disadvantaged minorities in science, technology, engineering, and mathematics (STEM) from grades K through 12.
- Increase the number of underserved and disadvantaged minorities that pursue PhD degrees and contribute to STEM related careers.

As a result, The Juxtopia Group, Inc. conducts empirical theoretical and application based research investigating effective learning technology tools and pedagogical approaches to improve, augment, and accelerate human learning performance. Through the Juxtopia Group, Dr. Doswell advises PhD students around the world with STEM concentrations and also mentors minority owned start-up businesses. He has personally mentored over 30 students of varying ages and from several countries ranging from India to Egypt and low income students in Baltimore, Virginia, and Washington, DC. Additionally, Dr. Doswell has addressed Congress on advanced technology areas.

Having served in leadership roles as consultant and/or trainer, Dr. Doswell, consulted with, and held leadership positions at several different companies and organizations, and federal agencies including Maryland Medical Systems, CompuServe, Lockheed Martin, BearingPoint, Scientific Applications International Corporation (SAIC), VirtualLogic, TRW, National Cancer Institute Center for Bioinformatics, among others. Dr. Doswell has spent over 13 years with a multidisciplinary concentration in advanced learning technologies and health informatics and health information technology research.

Dr. Doswell has also been quoted and features in national and local newspapers, magazines, TV, and radio including the Maryland Daily Record, Baltimore Washington Business Journal, Diverse Issues in Higher Education, Morgan State University radio station, MHz TV Networks, and CNN Headline news.

Dr. Doswell has also successfully won and executed grants from the National Science Foundation including a recent STTR grant with Georgetown University to create a novel Augmented Reality Intelligent Instruction System for use by manufacturers, medical professionals, US military war fighters, K-12 students, and to support homeland security first responders. Dr. Doswell also has experience writing and assisting other companies and universities write Small Business Innovative Research (SBIR) grant proposals for the National Institutes of Health, National Science Foundation, Department of Homeland Security, Department of Defense, and Department of Education.

Research Statement

Dr. Doswell's advanced learning technologies research focuses on learning technology interventions to improve human learning performance including mixed reality learning simulation, e-learning simulations, and video games for learning. He has developed and continually improves a unique system/software architecture for building families of virtual instructors or what Dr. Doswell has coined, the *Pedagogical Embodied Conversational Agent* (PECA). The PECA is an embodied artificially intelligent instructor that autonomously instructs human learners within mixed reality environments. To advance this multidisciplinary virtual instructor research, Dr. Doswell chairs the IEEE virtual instructor pilot research group that empanels international researchers for providing an international standard on building virtual instructors. He has also advised Gibbs College and Devry University on a novel video game curriculum for advancing science, technology, engineering, and math skills.

In health informatics and health information technology research, Dr. Doswell has lead research on fetal surveillance, telemedicine/telehealth, nanobiotechnology, and nanotechnology ethics. Dr. Doswell co-organized the American Public Health Association's first Health Informatics and Information Technology (HIIT) group and co-organized the first business meeting at APHA that discussed Nanotechnology for public health. Dr. Doswell also initiated and is leading an effort to engage minority serving institutions for health and learning technology research, which was endorsed by several members of the United States congress.

Dr. Doswell has presented his research to academic and non-academic conferences including, but not limited to the International Conference on Computer Information Technologies in Education, International Conference on Computer Graphics Imaging and Visualization, International Conference on Computer Graphics and Interactive Technologies, IEEE International Conference on Advanced Learning Technologies, Society of Applied Learning Technologies, International Digital Media Conference, International Museum and Web Conference, International Conference on Artificial Intelligence in Education, National Medical Association Annual conference, and American Public Health Association annual conference. He has also published industrial research in the IEEE Pervasive Computing magazine.

Dr. Juan E. Gilbert, Auburn University

<http://www.eng.auburn.edu/~gilbert>



Juan E. Gilbert was born in Hamilton, Ohio and was graduated from Hamilton High School. In 1991, he earned a Bachelor of Science degree in Systems Analysis from Miami University. Subsequently, he attained the Master of Science and Ph.D. in computer science from the University of Cincinnati in 1995 and 2000, respectively. In 2000, Dr. Gilbert joined the Computer Science & Software Engineering Department at Auburn University where he is currently the T-SYS Distinguished Associate Professor and directs the Human-Centered Computing(HCC) Lab, <http://www.HumanCenteredComputing.org/>.

Dr. Gilbert has research projects in Spoken Language Systems, Advanced Learning Technologies, User Interfaces (Usability), Ethnocomputing (Culturally Relevant Computing) and Databases. He has published more than 50 articles, given more than 100 talks and obtained more than \$3 million dollars in research funding in his seven years at Auburn University. Dr. Gilbert is a Senior Member of the IEEE Computer Society; he serves on the IEEE Computer Society Board of Governors; and he is the column Editor for the Broadening Participation in Computing Series in IEEE Computer magazine. In 2006 Dr. Gilbert was named one of the nation's top African-American Scholars by *Diverse Issues in Higher Education* and he was the recipient of the Black Engineer of the Year Special Recognition Award. In that same year, Dr. Gilbert also received the American Society for Engineering Education (ASEE) Minorities in Engineering Award. Recently, Dr. Gilbert was named a national role model by *Minority Access Inc.* At Auburn University, Dr. Gilbert has been honored with the Auburn University Alumni Engineering Council Junior Faculty Research Award, the Auburn University ACM Outstanding Faculty Member Award, the Auburn University Alumni Outstanding Minority Achievement Award and the Auburn University Outstanding Minority Service Award. In 2006, Dr. Gilbert was honored with a mural painting in New York City by *City Year New York*, a non-profit organization that unites a diverse group of 17 to 24 year-old young people for a year of full-time, rigorous community service, leadership development, and civic engagement.

Research Statement

Dr. Gilbert's research is in Human Centered Computing. The goal of this research is to design, implement and evaluate *innovative solutions to real world problems*. Human Centered Computing research exists between disciplines. Dr. Gilbert works on applied problems that require interdisciplinary solutions, specifically dealing with research issues in Natural Interactive Systems, Advanced Learning Technologies/Intelligent Tutoring Systems, Ethnocomputing and Information Technology Workforce, Human-Computer Interaction, Databases and Data Mining.

In Natural Interactive Systems (NIS), he is interested in creating pervasive user interfaces where the user interacts with the system using speech or multiple modalities. He is researching the design, implementation and the evaluation of naturally interactive systems. A good example of his

research in NIS is *Prime III*. Prime III is a secure, multimodal electronic voting system, <http://www.PrimeVotingSystem.com> and <http://www.youtube.com/watch?v=3egmhvgYil0>. Prime III provides an easy to use multimodal user interface that allows greater participation in the electoral process. Voters that can't read, hear, have visual impairments or physical impairments, can still vote using Prime III. In all of his speech enabled projects, he uses *Information Verbalization*, the use of computer supported, auditory interactions to amplify understanding of abstract and/or large data.

In Advanced Learning Technologies/Intelligent Tutoring Systems, Dr. Gilbert's research aims to create and study applications that employ intelligent strategies that personalize instruction. In some implementations, this involves the use of spoken language systems and Animated Pedagogical Agents (APAs). He researches the use and impact of culturally relevant environments that use culture in the education or training environment. This is a form of Ethnocomputing (<http://en.wikipedia.org/wiki/Ethnocomputing>). In his latest research efforts, he is studying game-like interfaces that provide naturally interactive instruction using animation, artificial intelligence, and speech. An example of this work can be seen at <http://www.aadmlss.com>.

In the area of Ethnocomputing or Culturally Relevant Computing, he is investigating the use and impact of culture in computing. This research suggests that culture can be used to increase interest, user satisfaction and ease of use in computing applications. In addition, Dr. Gilbert is working on information technology workforce issues. Specifically, he is investigating pedagogies and programs that broaden participation in computing for people in underrepresented groups. I am studying effective practices that help recruit, retain and graduate people from underrepresented groups in Science, Technology, Engineering and Mathematics (STEM).

In the area of Databases and Data Mining, Dr. Gilbert is investigating data mining algorithms and tools that answer complex questions from business intelligence, education, and society in general. For example, he uses clustering algorithms in a tool is called *Applications Quest* to process college admissions applications in order to increase holistic diversity, <http://www.ApplicationsQuest.com>

Dr. Shaun Gittens, Auburn University

<http://www.auburn.edu/~sdg0003/>



Dr. Gittens is a postdoctoral fellow at Auburn University located in Auburn, AL. His primary research interests include Machine Learning, Neural Computation, Computational Brain Modeling, Information Visualization, and Visual Data Mining.

Research Statement

My research interests focus on two differing areas of study: machine learning and information visualization. Within machine learning, my research primarily deals with improving supervised training methods of neural models that operate in external environments. My particular contributions in this area can have important implications in the development of improved learning capabilities for robots that aim to carry out sequential tasks with minimal feedback from the physical world in which they operate. I demonstrate the usage of this same work in my construction of a small, yet ambitious, attempt at simulating phoneme sequence acquisition in the brain consisting of four separate interacting computational neural models. On another note, within the area of information visualization, my research investigates methods that can be used to better visually inspect the results of various machine learning and data mining methods. These visual data mining methods become particularly valuable when used to analyze derived predictive models that become overly complex as well as when the data used to create such models are large in size.

Neural Network Sequence Generation in Sequential Environments

Consider, for instance, a task where a sightless robot, having no access to a reliable stream of camera images, attempts to “learn” to kick a soccer ball in a particular trajectory. This scenario would require a neural network learner to accept, as a training pair, the single intent to kick the ball in some unique fashion as input, along with the intended trajectory for the ball to traverse as output. Ultimately, the objective of the learner would be to learn the appropriate sequence of robot leg joint angles it needs to generate in order to achieve the desired results in kicking the ball. This would be considered a distal sequence generation learning task, where the intent to kick the ball represents the sole input, the camera images, or lack thereof, comprise the current state feedback from the environment, and the desired ball trajectory represents the sequence of distal desired outcomes sought in the environment. The sequence of kicking motions to be discovered and acquired by the distal learner comprises the proximal sequence of agent actions. Given some unique static input stimulus and in the likely absence of current state feedback from the environment, the goal of the distal sequential learner is to output the appropriate proximal action sequence which will ultimately yield the corresponding sequence of desired distal outcomes.

My dissertation research involved incorporating recurrent, or self-looping, structures accompanied by mechanisms which function as “memory” into the existing non-sequential distal supervised learning framework to better address sequential learning tasks designed to take place in an environment. Given the potential absence of current state feedback to the distal learner, the modified framework is designed to induce effective training of a recurrent neural network operating in an external environment. The effectiveness of this modified distal learning architecture is demonstrated in its application in the development of a simplified simulation of cognitive function acquisition which seeks to emulate the process that parts of the brain undergo when attempting to learn to produce phoneme sequences.

In order to construct such an ambitious environment, two important, yet complex, computational mechanisms were required to approximate the process of phoneme sequence generation: 1) a smooth mapping to characterize the disconnect between speaking (the mouth) and hearing (the ears), and 2) some mechanism for characterizing memory storage in the brain. A method was derived which was capable of transforming a finite look-up table of phoneme features into one smooth and continuous mapping from spoken utterances to heard sounds. Then, a slightly modified version of a sequential self-organizing map (specifically, the SARDNET architecture developed by James and Miikkulainen) was employed to function as associative memory in the brain for this particular task. Even provided the nature of this very complex environment, initial simulations of this cognitive function acquisition task were demonstrated to yield successful acquisition of up to 12 out of 15 desired words by the distal recurrent learner.

Future Research Directions

My future research plans regarding this work entail making this connectionist simulation of phoneme sequence acquisition more biologically plausible. Mechanisms observed in the human brain (e.g. increased interconnectivity among neighboring neurons and lateralization of the brain) can be incorporated into the existing model for improved study of the actual phoneme sequence acquisition process. Future attempts will be made to similarly replicate motor skill acquisition in the brain through use of mobile robots purposely situated in the physical world. Most computational simulations of brain behavior are developed on stationary computers that are devoid of actual real-world state feedback and incapable of influencing changes in their physical environment. The use of an architecture such as this which is capable of developing some measure of cognitive function acquisition in robots acting in the real world can potentially mark novel advancements in the development of neural models for computational brain study.

Structural Evaluation of Decision Trees Using Treemap Visualization Techniques

Research Overview

Currently, decision trees have been utilized in the creation of new insights from all types of classification data in a myriad of application areas ranging from engineering and business, to health and social issues. However, methods for visually evaluating the validity and effectiveness of decision trees in their entirety remain quite limited. For small decision trees having heights of four or fewer levels, the existing use of node-edge graph visualizations can suffice in many instances. However, the vast majority of decision trees, as a result of the complex and sizeable datasets they are used to characterize, tend to be incredibly large and, often times, too large to fit on a screen and be visually comprehensible.

My current research investigates the benefits of treemap visualization methods, originally introduced by Ben Shneiderman to view hierarchical and organizational data, for analyzing such complex decision trees. Treemap visualization is an alternative method of displaying a hierarchy of data within a constrained rectangular region on a screen using some series of partitioning such that all leaf nodes, each of which represents a data point or row, remain visible regardless of the number. Using the aforementioned treemap visualization technique, decision trees, after being reconstructed as equivalent classification trees over the training data, can be viewed in their entirety based on the space constraining properties typical in treemap visualizations. Subsequently, very important properties of the decision tree can be observed which otherwise could never be seen using the standard node-edge graph format. As a result of being able to view all the training data simultaneously within the structure of the decision tree, factors such as “bad” subtrees, outliers in the data, and evidence of overtraining can be immediately observed and investigated. In addition, this visualization technique will offer a more robust evaluation of differing decision tree methods and practices on the same dataset rather than merely comparing accuracy rates.

Future Research Directions

Subsequent research efforts in this area would entail utilizing treemap visualization for inspection of other types of hierarchical data mining methods (e.g. regression trees, cluster trees, etc.) and analysis of various current and up-and-coming decision tree algorithms and training practices. Ultimately, I envision developing new methodologies for constructing decision trees through interactive use of treemap visualizations either in an entirely manual manner or in conjunction with existing decision tree algorithms.

Dr. Illya V. Hicks, Rice University

<http://www.caam.rice.edu/~ivhicks>



Illya V. Hicks was born and raised in Waco, TX. Dr. Hicks lettered in football and earned a BS in mathematics (1995) from Southwest Texas State University (currently Texas State University at San Marcos). Dr. Hicks also received the MA and PhD in Computational and Applied Mathematics (2000) from Rice University while being supported under an AT&T Labs Fellowship. Dr. Hicks was a member of the Industrial and Systems Engineering Department faculty at Texas A&M University from 2000 to 2006. In 2006 Dr. Hicks moved to Rice University where he is currently an associate professor in the Computational and Applied Mathematics

Department.

Dr. Hicks' research interests are in combinatorial optimization, graph theory, and integer programming. Some applications of interest are in network design, cancer treatment, social networks, and logistics. Dr. Hicks was awarded the 2005 Optimization Prize for Young Researchers from the Optimization Society of the Institute for Operations Research and the Management Sciences (INFORMS).

Dr. Hicks is a cluster leader in the Rice/Houston Alliance for Graduate Education and the Professoriate (AGEP) program funded by the National Science Foundation and a past president of the Minority Issues Forum of INFORMS. In addition, Dr. Hicks served as the co-advisor for the National Society of Black Engineers (NSBE) undergraduate chapter and assistant director of the Sloan Foundation PhD Program while at Texas A&M University.

Research Statement

Numerous NP-complete problems become trivial on trees. For example, the Hamiltonian path problem and the chromatic number problem are both trivial on trees. For a particular problem modeled on a graph, one could work the problem on a proper subset of the edges and get a partial solution to the problem. Since this approach has a divide and conquer appeal, one could use a tree to represent some proper subsets of the edges where every leaf corresponds to an edge of the graph. Since any edge in the tree represents two edge-disjoint subgraphs and partial solutions, the interaction between the nodes that are in both subgraphs dictate the union of the two partial solutions into a solution toward the problem. The tree that decomposes the graph into edge-disjoint subgraphs is a foundational idea behind the usage of branch decompositions for dynamic programming algorithms for solving NP-hard problems (branch-decomposition-based algorithms). These algorithms are also part of the class of algorithms deemed fixed-parameter tractable algorithms.

Let G be a graph with node set $V(G)$ and edge set $E(G)$. Let T be a tree having $|E(G)|$ leaves in which every non-leaf node has degree 3. Let ν be a bijection (one-to-one and onto function) from the edges of G to the leaves of T . The pair (T, ν) is called a *branch decomposition* of G . Notice that removing an edge, say e , of T partitions the leaves of T and the edges of G into two subsets A_e and

B_e . The *middle set* of e and of (A_e, B_e) , denoted by $mid(e)$ or $mid(A_e, B_e)$, is the set $V(G[A_e]) \cap V(G[B_e])$ where $G[A_e]$ is the subgraph of G induced by A_e and similarly for $G[B_e]$. The *width* of a branch decomposition (T, v) is the maximum order of the middle sets over all edges in T . The *branchwidth* of G , denoted by $\beta(G)$, is the minimum width over all branch decompositions of G . A branch decomposition of G is *optimal* if its width is equal to the branchwidth of G . For example, Figure 1 gives an optimal branch decomposition of an example graph where some of the middle sets for the branch decomposition are provided.




Figure 1: Example Graph G with Optimal Branch Decomposition (T, v) with width 4

I have carried out a large body of work in branch decomposition research. The notions of branch decompositions and branchwidth were introduced by Robertson and Seymour to assist in proving the Graph Minors Theorem. My most significant contributions in this area include a practical implementation of a branch-decomposition-based algorithm for graph minor containment and practical implementations of algorithms to compute the branchwidth and an optimal branch decomposition of planar graphs. Also, I have introduced a new competitive heuristic algorithm for branch decompositions and an innovative branch-decomposition-based algorithm to compute an optimal branch decomposition for general graphs, which received the “Optimization Prize for Young Researchers” from the optimization society of INFORMS. The aforementioned results illustrate that branch decomposition techniques can be developed for practical algorithms to solve computationally hard problems. Furthermore, a colleague at Texas State University at San Marcos (my alma mater) and I have proved a conjecture stating that if a graph has a cycle of length at least two then the graph’s branchwidth is equivalent to the branchwidth of the graph’s cycle matroid. Most of this work was supported through NSF-funded grants.

Additionally, I have conducted research in independence systems. An independence system is a pair (E, \mathcal{I}) where E is a set of elements and \mathcal{I} is a family of subsets of E deemed “independent”. In addition, the pair (E, \mathcal{I}) have to satisfy two axioms: the empty set is “independent” and that the subset of an “independent” set is also “independent”. Many problems that are NP-hard can be modeled as an independence system such as the maximum clique problem, the maximum independent set problem, and the maximum planar subgraph problem. In reference to the latter problem, one of my papers found nine new facets for the planar subgraph polytope. The maximum planar subgraph problem has applications in network visualization and facility layout design. In conjunction with Jeff Warren, a past doctoral student, and others, I have implemented a new branch-and-price approach for the maximum weight independent set problem that performed well

for sparse graphs. Also, my student Ben McClosky and I looked at the composition of facets for the independent set polytope. The maximum weight independent set problem has applications in pattern recognition, computation biology, and combinatorial auctions.

Currently, I am working on the maximum k -plex (co- k -plex) problem. The k -plex problem is analogous to the maximum clique (stable set) problem in a more general case. A 1-plex in a graph is a clique. Social network analysts have been trying to characterize cohesive subgroups in social networks since the 1950's and invented these terms to describe a comprehensive structure of cohesive subgroups. A k -plex incorporates cohesiveness, reachability (fast communication among group members), robustness, and is adaptable to sparse networks. These degree-based relaxations of cliques can encompass more cohesiveness even when adequate information about relationships is missing (missing edges). In this regard, my research with my student Ben McClosky has derived new facets for the associated polytopes, extending known combinatorial algorithms for cliques to the k -plex problem, studied the co- k -plex polynomial, and made some tremendous strides for extending the perfect graph theorem to the k -plex case. In addition, constructing k -plexes can be vital for hierarchical infrastructures for ad-hoc wireless networks and data mining as well as social network analysis in terms of marketing and homeland security.

Dr. Charles Isbell, Georgia Institute of Technology

<http://www.cc.gatech.edu/~isbell>



Dr. Charles Lee Isbell, Jr. received his Bachelor of Science degree in computer science in 1990 from the Georgia Institute of Technology, where he was named its outstanding student by the President. Awarded a fellowship from AT&T Bell Labs as well as an NSF fellowship, he continued his education at the Artificial Intelligence Laboratory at the Massachusetts Institute of Technology. After earning his PhD from MIT in 1998, Charles joined AT&T Labs/Research. In the fall of 2002, he returned to Georgia Tech to join the faculty of the College of Computing.

Charles' research interests are varied. He has worked on a number of projects, including scaling machine learning algorithms like independent components analysis to problem spaces existing in hundreds of thousands of dimensions, developing extensions to description logics, developing new reinforcement learning techniques for balancing multiple sources of reward in social environments, state and activity discovery, light-weight coordination between multiple agents, partial programming, and building adaptive email architectures for users who need to manage hundreds of email messages a day.

The unifying theme of his work in recent years has been using statistical machine learning to enable autonomous agents to engage in life-long learning when in the presence of thousands of other intelligent agents, including humans. His recent technical focus is on developing new algorithms for activity discovery; doing adaptive coordination, especially in narrative; and in developing adaptive programming languages. His work with agents who interact in social communities has been featured in The New York Times, the Washington Post and Time magazine's inaugural edition of Time Digital magazine, as well as in several technical collections. Since graduating from MIT, he has won two best paper awards for technical contributions in this area. He was also awarded the Outstanding Junior Faculty Award in the College for some of this work.

Since returning to Georgia Tech, Charles has also pursued reform in computer science education. He has been awarded the William A. "Gus" Baird Faculty Teacher Award, and has been granted the Dean's Award for singular contribution to the College. The latter was for his work on Threads; Georgia Tech's new structuring principle for computer science curricula. This work has received international attention, and been presented in the academic and popular press.

Charles' research group is The Laboratory for Interactive Artificial Intelligence. For more information, please see: <http://www.cc.gatech.edu/~isbell/>.

Research Statement

My fundamental research goal is to understand how to build autonomous agents that must live and interact with large numbers of other intelligent agents, some of whom may be human. Progress towards this goal means that we can build usable artificial systems that can: work with humans to accomplish tasks more effectively; be more robust to changes in environment, relationships, and goals; and better co-exist with humans as long-lived partners.

In typical work on multi-agent systems, researchers deal with cases where the agents in question—be they physical robots or virtual agents—are working towards a common goal such as winning a soccer game or efficiently exploring terrain on Mars. By contrast, my work focuses on problems where this characterization does not hold: the agents in question are not only heterogeneous entities with heterogeneous abilities; they have heterogeneous and non-stationary goals. Further, because some of the agents may in fact be human, it is not possible to communicate intentions directly. The artificial agent cannot simply send an XML stream to the human and the human is not always capable of or interested in articulating her goals directly to the artificial agent. There are many interesting problems that have these properties, including:

- Distributed, personalized and adaptive training
- Hands-on interactive education
- The simulation and analysis of large social, political and economic systems
- Interactive entertainment such as massively multi-player games and interactive fiction

My contributions to this area have been a mix of foundational and engineering results. I have focused in recent years on how to best use statistical machine learning to address problems in domains that have the properties and constraints I have described above. On the one hand, this brings the tools of machine learning to these domains. On the other, making progress requires developing machine learning tools that work under circumstances that violate the usual assumptions (e.g., the environment is Markovian and stationary, there is a single source of reward, and so on).

My work is at the intersection of artificial intelligence, human-computer interaction and systems engineering. I would define my field of study as interactive artificial intelligence, and suggest that the central technical issues are:

- Adaptive modeling, including state and activity discovery
- Scalable coordination
- Development environments and methods that support the rapid prototyping

Adaptive Modeling

There has been a great deal of past work in developing algorithms and techniques that can recognize pre-defined activities. Although powerful, these techniques suffer from two problems. First, the patterns must be precisely specified. Second, and perhaps more importantly, by specifying specific patterns, the recognizer is unable to find patterns that the designer did not already know.

By contrast, my work is focused on the problem of discovering new patterns, particularly in a way that allows humans to inject feedback into the process with little extra effort. In order to do so well, I have to answer a basic question: when is a new pattern worth discovering? Or to put it another way, why is one way of looking at the world better than another? My answer is that the best indication that an activity is worth knowing is that detecting that activity empirically and demonstrably helps an agent to perform its task. My primary contributions in this area include: new algorithms and theory in constructivist learning; algorithms for prediction-based task discovery and supervised clustering; and extensions of reinforcement learning to social and human domains.

Scalable Coordination

Using models of itself and other agents, an agent should be able to identify potential mutually beneficial partnerships, and then act upon them. In focusing on collaboration, I seek to answer a set of questions: How can an agent detect that another agent can help it to accomplish its own goals? How can the agents build a representation of their shared goals that does not also require a heavy, centralized super mind? I have focused mainly on the latter question. My answer is to have the agent use models to detect the likely goals and plans of others and to use that to construct a lightweight manager that intervenes when the group is deviating from its shared goals. My primary contributions in this area include: establishing a variety of techniques for adaptive drama management; and the invention of Targeted Trajectory Distribution Markov Decision Processes (TTD-MDPs).

Development Tools

As a practical matter, it should be possible for human developers to program new agents and build new environments rapidly despite their complexity. In focusing on execution environments and executable models, we seek to answer questions revolving around systems engineering: How is it possible to allow non-machine learning experts to take advantage of machine learning algorithms in order to construct autonomous, adaptive systems? My answer is to integrate machine learning algorithms directly into programming languages and to use agent models that can be represented as operationalizable and executable code. What does this mean for building adaptive and interactive agents? What implications does such an idea have for a compiler? How does this affect the programmer in organizing and thinking about problems? In debugging? My primary contributions in this area include: developing an adaptive programming language called A²BL; and creating new tools for allowing non-experts to do rapid authoring of large-scale multi-agent games and simulations.

Machine Learning and Interactive Artificial Intelligence

All of my efforts have followed a cycle of identifying a domain or class of problems involving machines that must interact with and adapt to humans; articulating a manageable technical problem within that space; adapting or creating machine learning techniques to address them; and attempting to develop tools for making that process easier. Over time, I intend to bring the threads of my work further together, building end-to-end systems that exercise and integrate each stage of the pipeline, from authoring and development, to deployment, to long-term evaluation.

Dr. Odest Chadwicke Jenkins, Brown University

<http://www.cs.brown.edu/~cjenkins/>



Odest Chadwicke Jenkins, Ph.D., is an Assistant Professor of Computer Science at Brown University. He began this position in 2004 after serving as a postdoctoral researcher at the Center for Robotics and Embedded Systems at the University of Southern California. Prof. Jenkins earned his B.S. in Computer Science and Mathematics at Alma College (1996), M.S. in Computer Science at Georgia Tech (1998), and Ph.D. in Computer Science at the University of Southern California (2003). In 2007, he received Young Investigator funding from the Office of Naval Research and the Presidential Early Career Award for Scientists and Engineers for his work in learning primitive models of human motion for humanoid robot control and kinematic tracking. Prof. Jenkins' research group, "R-LAB: Robotics, Learning and Autonomy at Brown", explores research into realizing robots and autonomous systems as effective collaborators for pursuing human endeavors. His primary research topics lie in human-robot interaction and robot learning, with a specific focus on robot learning from human demonstration. His work also addresses research problems in robot/computer perception, humanoid robotics, machine learning, autonomous control, dexterous manipulation, computer animation, and game development.

Personal Biography

I was born on January 9, 1974 in Washington, D.C. to Odest Charles and Nadine Francis Jenkins. I have a younger brother, Oren, and sister, Nadalynn. I am married to Sarah Jenkins and we have two children, Morgan and Wesley. My research addresses various problems in autonomous robotics, namely how to teach robots from demonstration. In 2007, I received the Presidential Early Career Award for Scientists and Engineers, the highest honor bestowed by the U.S. government on outstanding scientists and engineers beginning their independent careers.

My parents were both born and raised in Jefferson County, Texas, not too far from Houston. Despite growing up in the segregated south, both of my parents able to pursue and complete undergraduate and graduate degrees. My mother (the actual Dr. Jenkins) received Masters and Doctorate degrees in Education and has pursued a career in teaching and administration at various respected institutions of higher learning. My father got a Masters degree in Sociology and pursued a career in the U.S. Federal Bureau of Prisons. His career in the Bureau of Prisons was rather successful, leading to him becoming a Warden at notable correctional facilities across the country. This rise up the career ladder, however, meant a childhood of moving from one locale to another (I never lived anywhere for more than 4 years until my pursuing doctorate in Los Angeles). Up through my high school graduation in New Jersey, my friends tended to be transitional from moment to moment rather than close and meaningful friendships.

My life was transformed Christmas Day 1981 with a gift from my parents. This gift was an Atari 2600 video game system and the beginning of my path towards a career in computing. Although crude by today's standards, games such as Space Invaders, Asteroids, and Breakout blew my mind with graphics that were interactive and displayed on the same television that would see my favorite (but non-interactive) shows. It all seemed like a wonderful magic that

would play with me and I could not get enough. I stockpiled a large collection of games, spanning the entertaining challenge of Pitfall to the disappointing port of Pac-Man. I particularly loved the sports games, when living in more isolated areas where pick-up games were impossible to find.

2600 games being as limited as they are, I always thought I could design better games if I only knew how they were made. This led me to obsessively find out as much as I could about making games throughout my childhood. I typically learned through various magazines and books that offered program listings in the BASIC language that I could type in and play on my Apple IIe. I would modify these programs to make my name appear in various places and make my own programs, such as one to print out my future Christmas lists. Although my playing of games remained constant, my involvement with computer programming went through peaks and valleys. My knowledge was mostly informal from teaching myself and I did not have adequate resources (this was way before the popular Internet). While having an education is extremely valuable, classes at school focused so much on the classical “liberal arts” topics and failed to engage my interest as much as video games. To put it mildly, my academic performance in high school was less than inspired and I had the grades to show for it. Although I really liked computers and loved games, I did not see myself as someone with the ability or skills to pursue a career in game development.

It was not until I got to college, in a supportive environment for computing, that I realized I had what it took to pursue my passion. I chose Alma College because it was a small college with a teaching emphasis. My original intention was to become an accountant because it seemed doable, but I took an introductory computer science class on the off-chance I liked it. I did really well in the course and, more importantly, had a revelation. It turns out that my obsession was with not just games, but computing and innovation as a whole. My passion was developing new technology to enable people across society to achieve more than they thought possible and experience the world in new ways, the same way the Atari changed my life.

At the time, the new transformative technology was 3D computer graphics, now ubiquitous in video games and animation, and I went to Georgia Tech to learn about it through a Masters program in CS. During my Masters program, I was fortunate to work with Prof. Jessica Hodgins, a pioneer in physics-based animation by combining techniques from animation and robotics. Although my education and training in graphics was excellent and could try a career in games, I felt that a career in research would lead to the true transformative innovations that I wanted to fulfill in my life. Following a pointer from Prof. Hodgins, I enrolled at USC and established relationships with both graphics and robotics researchers. Inspired by the work of Prof. Maja Mataric, my eventual doctoral advisor, I chose to pursue a dissertation in humanoid robotics and autonomous robot learning. After a lot of paper reading, constructive discussions, robot hacking, paper writing, and travel, I finished my dissertation in 2003, leading to my move to Brown University as a faculty member in 2004.

A little more bio on being faculty

Being faculty at a university is the best job on the planet. It is very much like being an entrepreneur starting up a company, except the measure of success is in ideas and inspiration not dollars. The academic environment allows me to work on whatever I deem is important for the

future of society as long as I can generate ideas of importance and inspire others to find the value in my ideas. My research ideas revolve around autonomous robotics, which I think will be one of the next transformative technological frontiers in the upcoming decades. Similar to computer networks and 3D graphics during the 1980s, robotics is coming to the point “personal robots” are slowly and increasingly affecting the daily lives of people and is poised for major innovations, analogous to the web browser, Toy Story, and the Playstation. More importantly, inspiring and advancing the lives and careers of future generations is a central theme of my position. A lot of my job is listening to people about their interests and vision for the world and then creating situations where they are inspired to excel. Despite common conceptions, being an academic in computing is more about personal relationships than sitting alone in a room. Putting the right people together on projects of common interest makes all the difference.

Research Statement

My research fundamentally attempts to address the question “How should we program robots?” More specifically, how will robots become effective collaborators for pursuit of human endeavors? This question of programming has been (more or less) addressed for personal computers (PCs), which is now an essential collaboration tool throughout society. PCs have standard paradigms for manipulating and processing information, writing software, mass distribution, distributed interoperability over networks, and applying software and hardware upgrades. Programming robots shares much in common with programming PCs with one major exception, uncertainty. For the most part, PCs can deterministically control the state of the system, as the information stored in bits. As a result, software written by a developer on one PC will (more or less) run the same on any PC. Robots, on the other hand, must interact with real world environments where they are faced with overwhelming degrees of uncertainty in estimating and manipulating the state of the system. Because of this uncertainty, robot software written by someone for a specific robot in specific environments will not behave the same when moved to a user’s environment, repurposed for some new task, or ported to a different robot. Consequently, most robots are limited to highly-specialized to perform autonomously or require overly burdensome human supervision and intervention.

Robot learning is one part of the answer to the problem; the full answer may be a combination of human demonstration and guidance with learning mechanisms for statistical regression and reinforcement learning. Significant subproblems involve heterogeneous human-robot teamwork, object manipulation, and enhanced teleoperation and neural prosthetics.

Some key areas of future work

Robot learning from demonstration can be viewed as a regression problem (e.g. learning robot soccer from demonstration).

Learning primitive human motion models for humanoid robot control, gesture recognition, and human-robot teams (iRobot collaboration) is a key area of research involving new work in physics-based human tracking.

Dexterous manipulation for enhanced human capabilities through teleoperation (NASA Robonaut collaboration) and neural prosthetics (similar to Aggeliki’s work)

Probabilistic models for multi-robot control and for providing a unified mathematical framework for existing and potential algorithms

Dr. Russ Joseph, Northwestern University

<http://www.ece.northwestern.edu/~rjoseph/>



Russ Joseph is Assistant Professor of Electrical Engineering and Computer Science at Northwestern University. His research interests are in power and reliability aware computer architecture. Specifically his ongoing projects examine cooperative hardware/software techniques that adaptively manage microprocessors to respond to environmental conditions, phase-sensitive program characteristics, and hardware failure. Joseph has been a Northwestern faculty member since 2004. He is the recipient of a 2007 NSF CAREER Award. He earned his Ph.D. in Electrical Engineering from Princeton University in 2004 under the direction of Professor Margaret Martonosi. His undergraduate studies were at Carnegie Mellon University where he double majored in Electrical and Computer Engineering and Computer Science. He received a B.S. degree with University and College of Engineering honors from CMU in 1999.

Research Statement

In leading edge microprocessor design, physical technology constraints will increasingly limit computing performance. Physical phenomena such as power dissipation, thermal cooling, fabrication errors, and component wear-out have a detrimental impact on clock frequency, time to failure, and energy efficiency. These phenomena cannot be fully addressed at design time because they: (1) are time sensitive or random in nature, (2) have complex interdependencies, and (3) call for situational or application sensitive solutions.

The primary goal of my work is to discover novel architectures that offer appropriate levels of performance, robustness, and efficiency through runtime adaptation and tradeoff.

My research approach is to first understand the relationship between physical phenomena which we hope to mitigate and microarchitectural design characteristics which we can control. This often leads to the development of physically grounded models which are either empirically constructed from circuit level simulation or analytically derived from first principles. My research draws on these models to develop flexible architectures that can operate at different points on the power/performance/reliability continuum. I then explore management policies that engage the flexibility in processor and chart a course for system.

An important component of my work examines the roles of hardware and software in processor management and exploits opportunities where the boundaries between the two are blurred. Increasingly, management solutions will require levels of sophistication that are not practical through conventional hardware design principles alone. In these cases, my work examines low-level software support via firmware and hypervisors that can apply global optimizations and state-space searches that are not hardware feasible. In the long term, this will lead to processor designs that go beyond self-healing and reactionary response to environmental stimuli. Future autonomous architectures will need to make inferences about their own status and predict future

behavior to proactively manage their resources. The following research projects explore supporting mechanisms and management strategies that move towards this vision.

Microarchitectural Design for Variability

Parameter variation due to manufacturing errors and circuit sensitivity to environmental conditions will be an unavoidable consequence of technology scaling in future generations. The impact of random and systematic variation in physical factors such as gate length, carrier mobility, and interconnect spacing will have a profound impact on clock frequency, power behavior, and reliability.

My research group is among the first to develop architecture-level models for parameter variability and microarchitectural components that can adapt to parameter deviations. When appropriate, we abandon the traditional egalitarian view of microarchitectural resources. We instead engineer systems that can understand localized parameter variations and treat their resources appropriately. We have developed novel cache organizations that allow a processor to adapt to local gate length variations which produce leaky subsections of the cache. Our work examines intelligent ways to disable portions of the cache to conserve power. Prioritized Cache Ways (PCW) and Prioritized Cache Sets (PCS) are dynamic cache resizing strategies that apply physical variation profiles improve energy-efficiency in high associativity and low associativity caches . We have plans to extend these concepts and techniques to other processor resources.

One of the goals of this project is to develop instance specific diagnostic and management techniques for chip multiprocessors (CMPs) under extreme parameter variability. To enable this, we have been working with the research group of my Northwestern colleague Professor Yehea Ismail to develop microarchitectural variation models that are based on circuit-level simulation. With these models as a foundation, we are in the process of developing thread and resource management policies that tailor their behavior specifically to individual CMPs which have their own variation profiles. Some of our recent work has looked at simple thread to core resource mapping strategies. The guiding principle is to export some knowledge of the instance specific variability to management software. The software can then appropriately select which threads should be run on which cores. Our simulations show an average 12% reduction in leakage power using variation aware management. We are currently exploring how resource matching strategies can be further aided by cooperative dynamic voltage frequency scaling (DVFS). This work is supported by NSF grant CCF-0541337.

Co-Designed Support For Hard Failure Recovery

As process technology scales, both fabrication induced and in-operation hard faults will become more prevalent, limiting yield and effective product lifetime. While transistor density will offer some hopes for robustness through resource redundancy, the sheer number of possible failures will complicate management of these redundant resources. The revitalization of machine virtualization offers several opportunities for hard failure tolerance and recovery.

My research group is exploring co-designed support for failure recovery via reliability focused virtual machine monitors (VMMs). These light-weight modules can be designed simultaneously

with the microarchitecture and can adapt hardware configurations to hide hard faults. The abstraction layer of the virtual machine hides the implementation details and failures from conventional operating systems and application code, eliminating the need for code modification. A recent workshop paper describes how this approach can be used to salvage partially functional cores on a chip multiprocessor [HPCRI-2006]. In particular, we demonstrate how a processor can recover from faulty functional units through emulation. We are currently developing a prototype VMM that runs on a full system CMP simulator. This will serve for the basis of our continued research on hard-fault recovery.

A major planned thrust of this work will examine techniques for predicting and avoiding hard faults that appear in the field. We plan to apply the variational performance/power models described above to examine long-term degradation due to electronic and thermal stress. Over time processor components will degrade over time (e.g. tolerate lower maximum clock frequencies) until they no longer give suitable performance or worse – give incorrect results at any clock frequency. We will examine management software that can monitor this degradation and administer in-field self-test. Based on degradation trajectories, the management software will steer execution and enable recovery strategies to prolong device lifetime. This work is supported by NSF grant CCF 0541337 and NSF CAREER Award CCF 0644332.

Dr. Ronald Metoyer, Oregon State University

<http://web.engr.oregonstate.edu/~metoyer/>



Ronald Metoyer is an Assistant Professor in the School of Electrical Engineering and Computer Science at Oregon State University. He received his Ph.D. from the Georgia Institute of Technology where he worked in the Graphics, Visualization and Usability Center. His research explores techniques for manipulating motion capture data and for facilitating the creation of 3D content by end users. His goal is to empower domain experts to create compelling and interactive content for their domain specific needs. In 2002, he received an NSF CAREER Award for his work in "Understanding the Complexities of Animated Content".

Research Statement

Research to Date

Computer generated 2D and 3D content is quickly becoming a part of our everyday lives, with 3D scenes showing up in our shopping, playing, training, and education experiences. As such, many content consumers are finding that they would like to create their own 2D and 3D content. With the complexity of the desired 3D content comes complexity in the tools to create the content – and this complexity leads to a widening gap between those who create content and those who consume the content. The overall goal of my research is to bridge this gap so that any content consumer, or end user, can also become a content creator and develop 3D content for their particular domain-specific needs.

To this end, my research program has been directed at the problem of making it easy for *novice* end users to create *compelling* 2D and 3D content with a focus on *human motion and behavior*. Compelling is often defined as “believable” or “realistic” in the graphics and animation domain. Before arriving at OSU, I addressed the creation of compelling behaviors such as pedestrian avoidance and football defensive maneuvers based on user provided examples. The idea behind this early work was that typical end users can physically demonstrate what they want to see (given an appropriate interface), even if they don’t have the skills to program that behavior into the system. This train of thought has continued to drive my research here at OSU.

To date, a significant part of my research thrust has been the generation and control of compelling, real-time, human motion based on example data. The most direct interface for demonstrating examples of human motion is through a motion capture system. I have developed a motion graph based approach to generating controlled motion from example motion capture data. The motion graph is essentially unrolled and spatial information is embedded into the new graph (called a mobility map) to achieve two objectives: real time rates and spatial control. This 3D implementation led directly to an innovative 2D motion model that was recently published in the pedestrian and evacuation dynamics area.

Modeling realistic 2D pedestrian motion. My model is the first based on motion capture graphs and was demonstrated to produce many of the characteristics of existing pedestrian models (lane formation, zippering, turn taking at doors) such as the social forces model (which I have also implemented and published). However, this model is especially novel and exciting because it is based on captured data. This means I can produce realistic models for any person whom I can motion capture (able bodied, disabled, children, elderly, etc.). The generation of 2D motion models for heterogeneous populations is a previously ignored but very important area that will ultimately lead to more realistic models of crowd movement. These will be especially useful in the study of large-scale panic and emergency evacuation situations.

Keyframing techniques for novice users. In 2005, I presented an innovative approach for keyframe animation that builds on the idea of end-user examples. Keyframing with existing tools can be difficult for novice users (especially the timing aspects) and generally requires artistic training to produce compelling results. Rather than time an animation in the traditional manner (set a keyframe value at a particular point in time), my approach allows the user to specify keyframe values first, without regard to time, and to later apply timing by “acting out” the desired timing information with a mouse or stylus input device. This approach was presented at the ACM SIGGRAPH Symposium on Computer Animation in 2005 and is the subject of an accepted journal article currently in revision.

Visual manipulation of simulations. Most recently, I have conducted research in an area called simulation programming, in which end users create content by programming a simulation to produce the desired behavior. This builds on my early work on creating behaviors by visual manipulation of simulations, but I have taken a programming approach as opposed to a machine learning approach. To start, I explored tangible user interfaces for allowing an end user, such as a football coach, to manipulate simulations. More recently, I focused on the language aspects of a visual strategy programming tool for football coaches to create animated 2D and ultimately 3D content. Rather than write scripts or program code, the user visually describes the desired behavior through direct manipulation of domain specific symbols. Our visual language for specifying football play strategy includes novel mechanisms for specifying sequential and parallel “and” behaviors without the burden of the English language confusions. This work is also represented in a collaborative effort in which we present the application of the cognitive dimensions framework to the design process of the visual language. Although this language is geared toward sports and gaming content, such languages have applications in the military domain where 3D content is used to train men and women for military situations. Advances in end-user visual languages for 3D content creation will ultimately lead to the production of targeted training and education material by domain experts themselves, rather than computer scientists and artists working with them. Furthermore, as the game community attempts to push the development of games to the end user, the requirement of powerful end-user tools and programming environments will become extremely important.

Human motion and perception. I have also worked on solutions to several animation problems not related to example-based motion or simulation programming. I recently embarked upon a collaborative effort with Dr. Victor Zordan from the University of California, Riverside, in which we generated compelling human preparatory responses for anticipating impacts and filled a void in previous research. I have also recently completed a perception study to better understand a common error in animation, known as “foot slipping”, and identified thresholds under which this error is imperceptible for stationary as well as moving cameras. This work is only a small part of a fairly open area of animation research that looks at what errors one can get

away with in current interactive 3D environments and how to quantify these errors. Such perception metrics will lead to more sophisticated techniques for optimizing processor usage when rendering complex 3D scenes with animated characters.

My expertise and potential research contributions were recognized in 2003 by NSF with the Faculty Early Career Award (Understanding the Complexities of Animated Content) which funded much of the above research. Along with the CAREER award, I have been recognized on several occasions as a promising young scientist. In 2004, I was invited to UC Irvine to keynote a session at the California Alliance for Minority Participation (CAMP) Symposium entitled “Road to the Professorate”. In 2005 I was invited to attend the 17th Annual Frontiers of Science symposium sponsored by the U.S. National Academy of Sciences. The symposia are meant to “... bring together some the very best young scientists to discuss exciting advances and opportunities in their fields....”

Future Research Directions

Over the next 10 years, I intend to continue to focus on end-user content creation, with less emphasis on human motion and more emphasis on 3D interaction and behavior. I believe that visual programming environments for 3D content represent the next big step in 3D gaming, education, and training environments. The development of visual methods for specifying complex interactions and behaviors within a 3D environment will lead to an explosion of interactive 3D content created by and for the novice. I recently submitted a grant proposal (with Dr. Erwig) to explore the visual programming of 3D mathematics “storygames,” where math concepts are manipulatable, interactive objects within the 3D environment. A second proposal will explore new models for novices to interact with and visualize the attributes and capabilities of 3D animated agents in situated environments through direct manipulation.

In a related vein, I intend to study the 3D web services area, which is growing as quickly as 3D gaming. The two are beginning to merge as we see web services centered on 3D content manipulation, sharing, delivery, etc. It will not be long before the modeling and animation services that are currently provided in expensive desktop software packages become available to the everyday user via web services. These services will undoubtedly require sophisticated yet simple interfaces for novice 3D users. The combination of usable visual programming environments for 3D behaviors with web services for the creation and sharing of 3D models, animation, and behaviors will facilitate the building of large databases of 3D content that include not only geometric models, but also 3D widgets, characters, and behaviors. Such content and interfaces can empower math teachers, for example, to create compelling and powerful content for expressing interesting problems and solutions and ultimately help to improve STEM (Science, Technology, Engineering, and Math) education in the United States.

Finally, I plan to work on the important problem of the shortage of IT/CS workers, as noted by the National Science Foundation. One solution supported by the NSF is to tap into a segment of the population currently underrepresented in computing – women and minorities. The first step in solving this problem is to begin to think about how to change the perception of the computing field in the eyes of these communities. I intend to pursue this clearly difficult problem by first understanding the current perception of computing and more importantly, identifying the attributes that would make it a more attractive option for women and minorities.

Dr. Hakim Weatherspoon, Cornell University

<http://www.cs.cornell.edu/~hweather/>



Hakim Weatherspoon is a post-doctoral fellow at Cornell University. His work covers various aspects of information systems, distributed systems, network systems, and peer-to-peer systems with focus on fault-tolerance, reliability, security, and performance of Internet-scale systems. He received his PhD from University of California, Berkeley in Computer science. He also received a Bachelors of Science degree in computer engineering from the University of Washington.

Research Statement

My research interests are broad, covering various aspects of information, distributed, network, and peer-to-peer systems. In these areas, I particularly focus on fault-tolerance, reliability, security, and performance of Internet-scale systems with decentralized—autonomous, federated, multi-organizational, and cooperative—control. My work generally applies principled approaches to understanding a problem, then validates the understanding gained through analysis by design, implementation, and evaluation of a system.

To date, I have applied this technique to several areas in systems research, including distributed wide-area storage, Byzantine fault-tolerance [4], efficient replica maintenance and erasure codes, structured and unstructured network overlays, content distribution networks, enterprise-level energy consumption and remote mirroring capabilities. I give a brief overview of these projects below.

Contributions

Distributed Wide-Area On-line Archival Storage Systems. Digital information plays an increasingly mission-critical role in military systems and other enterprises, and this trend has important implications. We need data storage systems that can ensure the durability, integrity, and accessibility of digital data, and do so under potentially turbulent conditions. For example, in a large scale system, servers continuously fail all the time; data should remain durable despite constant failure. As a more extreme example, during military engagements, communication links may be disrupted, systems may be destroyed or temporarily incapacitated, and load surges may result in degraded responsiveness; however, the availability of digital data is vital.

Antiquity is a distributed storage system designed for these sorts of challenging environments. It maintains data securely, consistently, and with high availability in a dynamic wide-area environment. At the core of the system is a novel secure log structure that permits Antiquity to guarantee the integrity of stored data, even under extreme stress. Data is replicated on multiple servers in a manner that ensures that data can be retrieved later even when some replicas are inaccessible. Moreover, unlike prior fault-tolerant systems, the Antiquity fault-tolerance protocols can handle high levels of server churn, regenerating data on the fly when necessary to handle faults ranging from server outages to Byzantine (malicious) attacks.

Antiquity's design has been validated via experimental evaluation that combined both global and local testbeds to realistically emulate the kinds of scenarios described above. Antiquity has been running on 400+ PlanetLab servers for over two months storing nearly 20,000 logs totaling more than 84 GB of data.

Despite continuous server churn, the logs used to locate and retrieve files remain durable. Antiquity was developed in the context of OceanStore. In particular, a component of OceanStore was a primary replica implemented as a Byzantine agreement process. This primary replica serialized and cryptographically signed all updates. Given this total order of all updates, the question was how to durably store and maintain the order. Antiquity's implementation of the interface and structure of a secure log assisted in durably maintaining the order over time. When data is later read from Antiquity, the secure log and repair protocols ensure that data will be returned and that returned data is the same as stored.

Efficient Replica Maintenance and Erasure Codes. Wide-area storage systems typically maintain data durability via fault tolerance and repair algorithms. Fault tolerance ensures that data is not lost due to server failure. Replication is the canonical solution for data fault tolerance. The challenge is knowing how many replicas to create and where to store them. Fault tolerance alone, however, is not sufficient to prevent data loss as the last replica will eventually fail. Thus, repair is required to replace replicas lost to failure. The system must monitor and detect server failure and create replicas in response. The problem is that not all server failure results in loss of data and the system can be tricked into creating replicas unnecessarily. The challenge is knowing when to create replicas. Both fault tolerance and repair are required to prevent the last replica from being lost, hence, maintain data durability.

My work, individually and with others, makes several contributions in this space and yields many insights on how to efficiently maintain data durability. First, I have explored the parameterization space of fault tolerance algorithms and associated durability. This work showed that erasure-coding reduces the bandwidth costs to maintain a target level of durability when compared to replication. Alternatively, for the same storage overhead and bandwidth costs, erasure-coding can maintain a significantly higher level of durability than replication. Next, it showed durability is related to the distribution of failure bursts. Further, it showed that random placement is sufficient to increase durability via reduced repair time and avoiding many correlated failures.

My work showed how to reduce costs (number of replicas created) due to transient failures. It demonstrated a principled way to estimate the amount of extra replication required to reduce repair costs due to transient failures. In particular, when data is immutable (i.e. cannot change), it showed that the system can limit the number of unnecessary copies made due to transient failures by ensuring that recovered copies are integrated in place into the replica set. The result is that the system performs some extra work for each object early in its life, but over the long term creates new copies of the object only as fast as it suffers permanent failures.

Using these insights, we developed the Carbonite replication algorithm for keeping data durable at a low cost. A simulation of Carbonite is able to maintain 100% durability over the course of a

year on PlanetLab despite losing over a third of the servers due to permanent failure such as disk failure. Further, in response to transient failures, it creates less than half the number of replicas when compared to previous systems.

Gossip-based Unstructured Overlays. Gossip-based mechanisms are touted for their simplicity, limited resource usage, robustness to failures, and tunable system behavior. Gossip is the periodic pairwise exchange of bounded size messages between random servers in a system, where the state of one (or both) changes to reflect the state of the other. Example environments where gossip-based mechanisms perform well include wide-area networks with limited access link bandwidth, wireless sensor networks with limited battery life and communication range, and mobile ad hoc networks with limited connectivity. I refer to the above environments as challenged networks.

I have become increasingly interested in exploiting gossip-based mechanisms in my work: they offer robust solutions to problems that can be hard to solve in other ways, such as providing durable storage in a challenged network. In particular, gossiping allows for proactively creating new replicas at a constant low rate in anticipation of future failures. Moreover, in power-constrained environments such as wireless sensor networks and mobile ad hoc networks, predictable communication patterns, as ensured by gossiping, facilitate energy savings. In my work with gossip, up to the present, I have demonstrated how to tune the rate of gossip to achieve desirable durability levels (rate of data loss) and energy savings.

On the other hand, I am cautious not to apply gossip for purposes to which it is ill-suited; it is not a “one-size-fits-all” mechanism. In particular, outside of these challenged networks or others like them, gossip may not be a good mechanism if server failure is a rare occurrence or new replicas can readily be created since resources such as bandwidth or battery life are not limited.

Looking towards the near future, I hope to explore applications of gossip to such problems as the convergence between structured and unstructured network overlays and rescue networks (i.e. sensor networks that provide durable storage until rescuers arrive).

Content Distribution Networks. Over the last few years there has been increasing usage of content distribution networks (CDNs) that deliver large volume data objects such as video and software. For example, approximately 70% of Tier 1 ISP traffic measured in Europe was peer-to-peer traffic; yet, the considerable bandwidth consumption is not necessary to satisfy the download demand. The challenge for CDNs like BitTorrent that transfer large objects is to minimize the download time while reducing network bandwidth consumption.

I have contributed to the design of two separate CDN projects. First, ChunkCast is an anycast service that optimizes large content distribution. It uses a distributed locality-aware directory that supports an efficient query for large content. It improves the median downloading time by at least 32% compared to previous approaches and emulates multicast trees without any explicit coordination of peers.

Second, AntFarm distributes content via managed swarms. Managed swarms differ from traditional swarming protocols in that they permit a single administrative entity, typically the content owner, to control the behavior of the swarms. This level of control exerted by a

coordinator, in turn, can enable the swarms to operate much more efficiently than both client-server systems and existing swarming systems. The central problem faced by such a service is how to minimize the time for content distribution, given the distributor's seeding bandwidth, the number of swarms, the number of peers in each swarm, churn, and the peers' upload and download capacities. Under the assumption that peers are selfish and potentially Byzantine, we first provide a protocol that enables the content distributor to detect misbehaving hosts, gather information on swarm dynamics, and direct the peers' allotment of upload bandwidth. We then derive the optimal strategy by which the content distributor can allocate its own seeding bandwidth and direct the upload bandwidths of peers to achieve the lowest possible delivery times for content. Extensive simulations and a PlanetLab deployment show that the system can significantly outperform BitTorrent.

Future Research: Optimizing Enterprise-level Remote Mirroring and Energy Consumption

A global network of datacenters is emerging as an important distributed systems paradigm—commodity clusters running high-performance applications, connected via high-speed, high-capacity, networks across hundreds of milliseconds of network latency. Some key challenges in this environment are the tradeoffs between disaster tolerance and performance and energy consumption of datacenters. However challenging, this paradigm provides for many rich research opportunities into problems that have yet to be solved. Further, my skills and qualification position me as a key player in the area.

First, disaster tolerance and recovery designs confront a conundrum: the tradeoff between either the loss of performance or the potential loss of data. On the one hand, loss of performance such as application response time or throughput may not be an option. For example, it may not be desirable to slow application response time until it is assured that data will not be lost in the event of disaster. Similarly undesirable is the prospect of data loss, which can be catastrophic for many companies and organizations — 70% of small firms that experience a major loss of data go out of business within a year[Source: DTI/Price Waterhouse Cooper, 2004]. Unfortunately, there is not much of a middle ground in the design space and designers must choose one or the other.

I am beginning to explore new ideas that may offer a middle ground. One idea is to proactively inject redundancy into the network to prevent loss of transmitted packets, and expose the status of out going data, so that the sender can resume activity as soon as a desired level of in-flight redundancy has been achieved for any given packet. For example, an optical link that drops one out of every trillion bits or 125 million 1 Kb packets (this is the maximum error threshold beyond which current carrier-grade optical equipment shuts down) can be pushed into losing less than 1 out of every 10¹⁶ packets by the simple expedient of sending each packet twice — a figure that begins to approach disk reliability. Utilizing the extra bandwidth available to send more redundancy codes can effectively turn the network into a portable disk where data is stored and shipped across to the remote mirror. More importantly, the idea is independent of link latency and can run at very high absolute data rates. This may create a viable new option for developers of mission-critical datacenters on the outer edge of the performance curve.

Second, energy consumption in modern datacenters is an issue receiving increasing attention because of its growing importance — an article in The New York Times describes one Google’s datacenters: “... a computing center as big as two football fields, with twin cooling plants protruding four stories into the sky”. At issue is the financial and environmental cost of keeping hundreds of thousands of disks spinning.

One of many possible opportunities for saving energy in large datacenters is within the file system. The idea is elegant in its simplicity: log structured file systems write only to the log head; as a result, if read accesses are served by the cache, then write accesses touch only the log head disk, potentially allowing us to power down all the other disks. Existing solutions like disk management solutions and caching solutions are typically application-specific; our solution, on the other hand, is applicable to any cacheable dataset. Since existing solutions are typically layered on top of the file-system, they could be used in conjunction with our solution to take advantage of application-specific optimizations

Looking further into the future, I am starting work on the conceptual design a new kind of file system that would link datacenters over very long distance optical network links. This raises many technical challenges, some of which are described above. Once finished, the file system will give users a very simple normal view of one file system, but should be able to automatically move data around for performance reasons, and also adapt itself to power and other environmental considerations. It will capitalize on opportunities to save power such as trading off which disks are spinning in which datacenters against the costs of communicating over fast, wide-area links.