

Tenth BU/CS IAP Research Open House

March 19, 2010

Poster Abstracts

The following are the abstracts of posters and demos on display in the CS Research Lab. Abstracts are listed (alphabetically by student name) under the broad areas of:

- DB** Databases
- IVC** Image and Video Computing
- N** Networking
- OS** Operating Systems
- PL** Programming Languages
- S** Security
- T** Theory

Databases

DB-1 k-Nearest Neighbors in Uncertain Graphs (Poster)

Michalis Potamias

Supervised By: George Kollios

Complex networks, such as biological, social, and communication networks, often entail uncertainty, and thus, can be modeled as probabilistic graphs. Similar to the problem of similarity search in standard graphs, a fundamental problem for probabilistic graphs is to answer efficiently k-nearest neighbor queries (k-NN), which is the problem of computing the k closest nodes to some specific node. We introduce a framework for processing k-NN queries in probabilistic graphs. Within this framework, we propose novel distance functions that extend well-known graph concepts, such as shortest paths. In order to compute them in probabilistic graphs, we design algorithms based on sampling and novel graph-transformation ideas. During k-NN query processing we efficiently prune the search space using new techniques. Our experiments demonstrate that our distance functions outperform alternatives in identifying true neighbors in real-world data. We also demonstrate that our algorithms scale for graphs with tens of millions of edges.

Image and Video Computing

IVC-1 Layered Graphical Models for Tracking Partially-Occluded Objects (Poster)

Vitaly Ablavsky

Supervised By: Stan Sclaroff

We propose a representation for scenes containing relocatable objects that can cause partial occlusions of people in a camera's field of view. In many practical applications, relocatable objects tend to appear often; therefore, models for them can be learned off-line and stored in a database. We formulate an occluder-centric representation, called a graphical model layer, where a person's motion in the ground plane is defined as a first-order Markov process on activity zones, while image evidence is aggregated in 2D observation regions that are depth-ordered with respect to the occlusion mask of the relocatable object. We represent real-world scenes as a composition of depth-ordered, interacting graphical model layers, and account for image evidence in a way that handles mutual overlap of the observation regions and their occlusions by the relocatable objects. These layers interact: proximate ground plane zones of different model instances are linked to allow a person to move between the layers, and image evidence is shared between the observation regions of these models. We demonstrate our formulation in tracking pedestrians in the vicinity of parked vehicles. Our results compare favorably with a sprite-learning algorithm and with a pedestrian tracker based on deformable contours.

IVC-2 Learning Actions From the Web (Poster)

Nazli Ikizler Cinbis, R. Gokberk Cinbis

Supervised By: Stan Sclaroff

We propose a generic method for action recognition in uncontrolled videos. The idea is to use images collected from the Web to learn representations of actions and use this knowledge to automatically annotate actions in videos. Our approach is unsupervised in the sense that it requires no human intervention other than the text querying. Its benefits are two-fold: 1) we can improve retrieval of action images, and 2) we can collect a large generic database of action poses, which can then be used in tagging videos. We present experimental evidence that using action images collected from the Web, annotating actions is possible.

IVC-3 Tracking Using Active Hidden Models (Poster)

Samuel Epstein

Supervised By: Margrit Betke

We introduce Active Hidden Models (AHM) that utilize kernel methods traditionally associated with classification. We use AHMs to track deformable objects in video sequences by leveraging kernel projections. We introduce the Kernel Semi-least Squares problem, whose solution is the subset projection method which improves the efficiency of our tracking approach by a factor of five. We successfully tested our method on facial tracking in cluttered backgrounds with extreme head orientations and fast head movements, on nonlinear deformation of surfaces, and on facial expression recognition.

IVC-4 HAIL: Hierarchical Adaptive Interface Layout (Demo and Poster)

John Magee

Supervised By: Margrit Betke

We present an adaptive user interface for individuals with severe motion disabilities who use mouse substitution interfaces. Typically, users are required to adapt to the interfaces that they wish to use. Instead, we propose interfaces that should change and adapt to the user and their individual abilities. The Hierarchical Adaptive Interface Layout (HAIL) model is a set of specifications for the design of user interface applications that adapt to the user, rather than requiring the user to adapt to the interface. In HAIL applications, all of the interactive components take place on configurable toolbars along the edge of the screen. We show two HAIL-based applications: a general purpose web browser, and a twitter client.

IVC-5 A Probabilistic Framework to Exploit Phonological Constraints with Non-rigid Image Alignment for Singer Independent Handshape Inference in American Sign Language Video (Poster)

Ashwin Thangali

Supervised By: Carol Neidle, Stan Sclaroff

Handshape plays an important role in signed languages among other signing parameters such as hand location, orientation, movement trajectory and facial actions. Recognizing handshapes in isolated hand images from single-view signing video is an insufficiently constrained problem since many handshapes share similar 3D configurations and are indistinguishable for some hand orientations. Additionally, significant changes in hand appearance are induced by the articulated structure of the hand and variants produced by different signers. Phonological rules in American Sign Language (ASL) govern the occurrence of handshapes in a sign and can be exploited to impose constraints during recognition. Signs in ASL are categorized on the basis of different compositional principles, monomorphemic lexical signs (the focus of this study) form the largest class among signs that conform to phonological rules; handshapes in these signs are specified for the articulating hand(s) at the start and end of the sign. We propose a probabilistic framework to exploit phonological constraints that govern the allowable start-end handshape pairs; the constraints can broadly be classified as bilateral handshape congruence (or distinctiveness), start-end handshape co-occurrence relationships and allophonic handshape substitutions to include handshape variants. We derive an algorithm to estimate parameters for the proposed model using handshape annotations for isolated signs signed by native signers; the trained model is used to infer start-end handshape pairs in a query sign. In the second part of this study, we develop a non-rigid image alignment algorithm to obtain a similarity measure for pairs of hand images from ASL video with improved robustness to variability inherent in handshape appearance. The proposed algorithm solves a sequence of sparse linear system of equations to compute the non-rigid alignment. This yields an algorithm that is significantly faster than state-of-the-art non-rigid alignment techniques based on Loopy Belief Propagation (LBP). In a feasibility study we demonstrate that the proposed non-rigid image alignment improves handshape recognition accuracy on an ASL handshape dataset; the query and database hand images are obtained from video of male and

female native signers signing from an ASL dictionary. The recognition results using the proposed approach compare favorably with the benchmark LBP algorithm while yielding a significant speedup.

IVC-6 Fast Globally Optimal 2D Human Detection with Loopy Graph Models

Tai-Peng Tian

Supervised By: Stan Sclaroff

We presents an algorithm for recovering the globally optimal 2D human figure detection using a loopy graph model. This is computationally challenging because the time complexity scales exponentially in the size of the largest clique in the graph. The proposed algorithm uses Branch and Bound (BB) to search for the globally optimal solution. The algorithm converges rapidly in practice and this is due to a novel method for quickly computing tree based lower bounds. The key idea is to recycle the dynamic programming (DP) tables associated with the tree model to look up the tree based lower bound rather than recomputing the lower bound from scratch. This technique is further sped up using Range Minimum Query data structures to provide $O(1)$ cost for computing the lower bound for most iterations of the BB algorithm.

IVC-7 Reconstruction and Analysis of 3D Trajectories of Brazilian Free-tailed Bats in Flight (Poster)

Diane Theriault and Zheng Wu

Supervised By: Margrit Betke

We developed a multi-camera method to reconstruct the 3D trajectories of animals in flight that uses a multi-hypothesis tracking approach and stereoscopy. We applied the method to analyze the infrared thermal videos captured by three spatially calibrated and temporally synchronized cameras. The videos show a column of Brazilian free-trailed bats emerging from their day-time roosts and flying to their foraging areas. We analyzed 7,000 individual flight trajectories and the spatial structure of the column of trajectories. We report estimates of the velocities of bats in flight, their distances to each other, and the angles subtended by the bats and their nearest neighbors.

IVC-8 Camera Self-Calibration, Plane Segmentation and 3D Reconstruction from Image Sequences of Polyhedral Scenes (Demo and Poster)

Flavio Viguera

Supervised By: Stan Sclaroff

This work mainly addresses the following problems: (i) self-calibration of a moving camera observing a 3D scene composed by planar structures; (ii) the scene segmentation and reconstruction. We propose an iterative linear algorithm exploiting the constraints induced by planarity in the scene. Instead of solving a non-linear or a complex multi-linear problem, we solve iteratively several linear problems: coplanar features segmentation, planar projective transferring, epipole computation, all plane intersections and calibration. Linear methods allow our approach to be suitable for real-time localization and 3D reconstruction.

Networks

N-1 Pricing Algorithms For a Two-Sided Internet Advertisement Market – Work In Progress (Poster)

Joseph Akinwumi

Supervised By: John Byers, Evimaria Terzi

The Google AdSense Program is a successful internet advertisement program where Google places contextual advertisements on third party websites and shares the resulting revenue with the website publishers. Advertisers have budgets and bid on ad slots while publishers set reserve prices for the ad slots on their websites. Following previous modeling efforts, we model this system as a two-sided market and define fairness and optimality solution concepts. We show a reduction from the Generalized Assignment Problem (GAP) to the problem of computing revenue-maximizing allocation and prices of publisher slots in our two-sided market under a first-price auction model. GAP is APX-hard, but a $(1 - 1/e)$ -approximation algorithm is known. Our current work involves extending the model to capture incentive-compatibility via second-price auctions, as well as scenarios with online or incomplete information.

N-2 Hyperbolic Scaling of Network Graphs (Poster)

Andrej Cvetkovski

Supervised By: Mark Crovella

Motivated by the explosion in availability of graph data measured from real-world communication networks, and more recently, networks of social relationship, we constructed MDS-PD, an embedding and visualization algorithm for dimensionality reduction of large scale sparse datasets. MDS-PD is a projective algorithm designed specifically for the Poincare disk model of the hyperbolic plane that can be used on multidimensional datasets as well as on bare inter-object distances. Here we present examples of use of our algorithm that demonstrate its utility in providing advanced understanding and novel insights into fundamental properties of network graphs derived from real-world measured data. We believe this novel tool will prove useful across a wide range of subsequent network measurement inquiries and the insights gained from its application will result in new solutions to several current networking problems.

N-3 HEAD in the Cloud (Poster)

Flavio Esposito

Supervised By: Ibrahim Matta

Providers are approaching new business models where virtualization-oriented architectures enable new services. In such paradigm, resource discovery and allocation are two challenging problems due to their online nature, and to the vast amount of heterogeneous resources available. Inspired by the design mechanisms for wide scale virtualization-based testbeds (i.e: GENI, PlanetLab, VINI, Emulab), we propose a new problem: Hierarchical Expectation-based Allocation Discovery (HEAD) for cloud

computing architectures. We formulate a solution of the general HEAD problem applying dual decomposition, and we show that simultaneous optimization of resource discovery and allocation leads to a more efficient management of the virtualizable resources.

N-4 Universal Mobility Modeling for MANETs (Poster)

Gonca Gursun

Supervised By: Ibrahim Matta

In order to simulate a new protocol for a Mobile Ad Hoc Network (MANET), it is necessary to use a mobility model that thoroughly represents the movement pattern of MANET agents. The choice of mobility models has a great importance as they impact the performance of the protocols used on top. Therefore, it is crucial to develop a deep understanding of mobility models. The development of a new mobility model is approached by establishing a one-to-one correspondence between application scenarios of the protocols and the mobility patterns of nodes. Therefore, a new model should be developed for every new application. In order to overcome this problem and provide reusability, several general synthetic mobility models have been proposed. However, these generalized models have the following drawbacks: (1) they only capture simple scenarios under specific assumptions, therefore they can not be used for complex mobility patterns, (2) they are not realistic. Consequently, these synthetic models lead to faulty conclusions when used to simulate MANET protocols. In order to solve these issues, we propose a unifying mobility model framework that creates realistic mobility models and enables the specification of complex mobility behaviors. Our framework leverages the Artificial Intelligence (AI) Game Programming techniques to define MANET agents that act intelligently. These agents can adjust their movements according to the changes in the environment and the movement of other agents.

N-5 Colocation as a Service: Strategic and Operational Services for Cloud Colocation (Poster)

Vatche Ishakian, Raymond Sweha, Jorge Londoño

Supervised By: Azer Bestavros

By co-locating with other tenants of an Infrastructure as a Service (IaaS) offering, IaaS users could reap significant cost savings by judiciously sharing their use of the fixed-size instances offered by IaaS providers. This poster presents the blueprints of a Colocation as a Service (CaaS) framework. CaaS strategic services identify coalitions of self-interested users that would benefit from colocation on shared instances. CaaS operational services provide the information necessary for, and carry out the reconfigurations mandated by strategic services. CaaS could be incorporated into an IaaS offering by providers; it could be implemented as a value-added proposition by IaaS resellers; or it could be directly leveraged in a peer-to-peer fashion by IaaS users. To establish the practicality of such offerings, we present XCS – a prototype implementation of CaaS on top of the Xen hypervisor. XCS makes specific choices with respect to the various elements of the CaaS framework: it implements strategic services based on a game-theoretic formulation of colocation; it features novel concurrent migration heuristics which are shown to be efficient; and it offers monitoring and accounting services at both the hypervisor and VM layers. Extensive experimental results obtained by running PlanetLab

trace-driven workloads on the XCS prototype confirm the premise of CaaS – by demonstrating the efficiency and scalability of XCS, and by quantifying the potential cost savings accrued through the use of XCS.

N-6 Incentive Mechanisms for Scheduling and Load-Balancing Atomic and Fluid Tasks on a Congestible Resource (Poster)

Jorge Londoño

Supervised By: Azer Bestavros

Commonly used pricing schemes such as the flat-rate or proportional-to-usage models do not provide incentives for applications with delay-tolerant tasks to schedule execution during low-demand periods. As a result, these tasks contribute to increased utilization during periods of high demand, degrading the performance of the system, and creating an otherwise unnecessary need for over provisioning. This work presents a incentive mechanism so that rational, selfish users benefit by scheduling their delay-tolerant tasks in low-demand periods, thus helping load-balance the overall resource utilization.

N-7 Analytical Understanding of Peer-to-Peer File Sharing (Poster)

Michel Machado

Supervised By: John Byers

Most of the results around P2P networks are rather experimental. Beyond the loose idea that peers must cooperate and be fair to each other, the literature provides neither a utility function, nor an optimization problem that rational peers solve. The best justification for the use of tit-for-tat strategy is still the victory of this strategy in the computer tournament run by Robert Axelrod in 1979. Moreover, the few general, analytical results do not prescribe how peers should behave. Our work models P2P networks using (Modern) Portfolio Theory, derives a utility function for peers minimizing their download time, analytically justifies the use of tit-for-tat, prescribes behavior that allows a peer to locally identify free-riders, and asserts that at most a peer does not have its uplink fully utilized when the system reaches steady state.

N-8 Principles of Safe Policy Routing Dynamics (Poster)

Karim Mattar, Samuel Epstein

Supervised By: Ibrahim Matta

We introduce the Dynamic Policy Routing (DPR) model that captures the propagation of route updates under arbitrary changes in topology or path preferences. DPR introduces the notion of causation chains where the route flap at one node causes a flap at the next node along the chain. Using DPR, we model the Gao-Rexford (economic) guidelines that guarantee the safety (i.e., convergence) of policy routing. We establish three principles of safe policy routing dynamics. The non-interference principle provides insight into which ASes can directly induce route changes in one another. The single cycle principle and the multi-tiered cycle principle provide insight into how cycles of routing updates can manifest in any network. We develop INTERFERENCEBEAT, a distributed algorithm that propagates a small

token along causation chains to check adherence to these principles. To enhance the diagnosis power of INTERFERENCEBEAT, we model four violations of the Gao-Rexford guidelines (e.g., transiting between peers) and characterize the resulting dynamics.

N-9 Angels In the Cloud: An On-Demand Peer-Assisted Content Distribution Cloud Service (Demo and Poster)

Raymond Sweha, Vatche Isahakian

Supervised By: Azer Bestavros

This project develops a Cloud Service for Internet content distribution. Our system assists Seeders (content originators) with the dissemination of content (a file) to a set of nodes using Peer-to-Peer (P2P) concepts so that this dissemination is completed in the minimum time possible. Prior results of ours suggest that minimizing content distribution time may be achieved by adding nodes that are not themselves interested in downloading the content, but rather in assisting other nodes with their download in a prescribed (provably optimal) fashion. We call such nodes Angels. The emerging cloud computing architecture offers the best mechanism to allow such Angels to be created on-the-fly. As seeders request assistance with their file distribution, our service responds by spawning virtual machines that act as Angels. In this work, we describe the design and implementation of our "Angels-on-Demand" cloud service as well as the API for invoking this service.

N-10 AS-level Topology Generation Incorporating Economic Principles (Poster)

Chong Wang

Supervised By: John Byers

We consider the problem of generating a representative autonomous system (AS) level topology of the Internet. Traditional models for network formation using graph-theoretic or optimization-based approaches do not capture many of the practical considerations facing ASes. In particular, the process by which ASes evolve is the outcome of cost-benefit analysis performed by each AS when establishing or renewing service contracts with other ASes. Based on observations from publicly available data, and incorporating other recent measurement and modeling results, we explicitly model the costs and actions performed by ASes to design a new AS topology generation framework incorporating economic principles. In our model, each AS selfishly optimizes its time-evolving payoff subject to local economic constraints, and its locally optimized decisions have influence over the whole network. We formulate the payoff function and optimization problem for each AS, devise efficient mechanisms to solve the problem and relate various properties of our generated topologies to those from previous models and observed topologies.

Operating Systems

OS-1 A Type-Theoretic Framework for Efficient and Safe Colocation of Periodic Real-time Systems (Poster)

Vatche Ishakian

Supervised By: Azer Bestavros and Assaf Kfoury

Desirable application performance is typically guaranteed through the use of Service Level Agreements (SLAs) that specify fixed fractions of resource capacities that must be allocated for unencumbered use by the application. The mapping between what constitutes desirable performance and SLAs is not unique: multiple SLA expressions might be functionally equivalent. Having the flexibility to transform SLAs from one form to another in a manner that is provably safe would enable hosting solutions to achieve significant efficiencies. This work demonstrates the promise of such an approach by proposing a type-theoretic framework for the representation and safe transformation of SLAs. Based on that framework, it describes a methodical approach for the inference of efficient and safe mappings of periodic, real-time tasks to the physical and virtual hosts that constitute a hierarchical scheduler. Extensive experimental results support the conclusion that the flexibility afforded by safe SLA transformations has the potential to yield significant savings

OS-2 A Real-Time USB subsystem for Quest Operating System (Poster)

Ye Li

Supervised By: Rich West

Operating Systems in Real-Time/Embedded systems must address task timing requests. However, existing OSes, e.g. RTLinux, RTAI, VxWorks, QNX, etc. still have areas of timing vulnerabilities such as shared caches, memory bus bandwidth contention and interrupts scheduling. Our previous work on processor aware interrupt scheduling and accounting partly addressed the problem of mismatched task and interrupt scheduling in modern systems. To add temporal isolation, we are now building a VCPU architecture for I/O and task scheduling as a Real-Time I/O subsystem in the Quest Operating System. USB is one of the most flexible, versatile and widely supported I/O subsystems used today. Consequently, our Quest I/O subsystem is focusing on predictable I/O management of USB devices. This poster describes early developments in the design of our USB-based real-time I/O subsystem, which we hope to use as part of our RacerX project.

OS-3 Virtual-CPU Architecture for Predictable Computing (Poster)

Meha Mehta

Supervised By: Richard West

Existing real-time systems still suffer from unpredictability due to issues like interrupts from I/O devices, shared caches in multiprocessor systems, priority inversion, deadlocks and memory bus contention. Our current work focuses on unpredictability due to interrupts from I/O devices. We propose a new "Virtual CPU" architecture that is based on bandwidth preserving servers like Sporadic

Servers and Constant Utilization Servers. VCPUs can be associated with conventional tasks as well as asynchronous event handlers. We implement this architecture in Linux Kernel and empirically show improvements in Linux Kernel predictability.

OS-4 RacerX: High Speed, Autonomous Buggy (Poster)

Eric Missimer, George Pabst

Supervised By: Richard West

RacerX encompasses several endeavors designed to push the limits of real-time systems. The RacerX buggy has been designed using off-the-shelf, commodity hardware and is currently running a popular Linux distribution. Plans are in the works to convert to the Quest research operating system for better real-time performance. We are currently working on having our RacerX navigate autonomously around a track. Our research has been focused this semester on developing fast algorithms for processing video frames from the on-board camera to identify the track and to send the appropriate RF signals to control the engine.

OS-5 Programmable Smart Machines: Exploiting Structure in Computation

(Poster)

George Pabst, Andrew Violino

Supervised By: Jonathan Appavoo

Our work this semester focuses on creating an execution signal encompassing the entire operation of a modern computer. The significant number of instructions and interrupts in a computer's operation make recording all of the events a challenging task. We are developing efficient algorithms and structures to record the execution and state of a machine while still providing a responsive user experience. Based on our early observations that there are visible patterns in the execution signal, we believe it would be possible to develop a mechanism which learns from and automatically optimizes the execution based only on the observed structure. We are exploring the use of Hierarchical Temporal Memory for this purpose.

Programming Languages

PL-1 OS Development with ATS (Poster)

Matthew Danish

Supervised By: Hongwei Xi

Typical operating system design is marked by trade-offs between speed and reliability, features and security. Most systems are written in a low-level untyped programming language to achieve optimal hardware usage and for other practical reasons. But, this often results in CPU, memory, and I/O protection flaws due to mistakes in unverified code. On the other hand, fully verified systems are exceedingly hard to construct on any industrial scale. A high-level programming language, with an expressive type system suitable for systems programming, can help alleviate many of these problems without requiring the enormous effort of full verification.

PL-2 An Automated Assistant for Formal and Semi-formal Reasoning (Demo and Poster)

Andrei Lapets, Prakash Lalwani, David House

Supervised By: Assaf Kfoury

Machine verification of formal arguments can only increase our confidence in the correctness of those arguments. However, the costs of employing machine verification still outweigh the benefits in a variety of formal reasoning activities. While it is by restricting a user to correct arguments that a machine verifier serves its purpose, such restrictions can inhibit even an expert user's productivity when they are reflected in the machine verifier's interface. To moderate these conflicting design constraints, we advocate an approach that separates a machine verifier into two layers: validation and interface. The validation layer is exclusively responsible for corroborating that a formal argument is sound with respect to a domain model. The interface layer is exclusively responsible for receiving and comprehending a formal argument, which may involve simulating the context (mental domain model and thought process) of the argument's human author. To evaluate the feasibility and effectiveness of a distinct interface layer for machine verification, we have developed aartifact. The aartifact system provides a familiar and flexible syntax that supports user-defined predicates and operators, a basic automated logical inference procedure, and a simulated context (static and dynamic). The simulated static context is an extensive library (implemented as a relational database) of thousands of common mathematical concepts and propositions indexed by their structure. The simulated dynamic context is a data structure (a transitively closed hypergraph) that mirrors a user's mental model of a formal argument. The dynamic context is derived automatically from the static context, and can recognize a user's implicit invocation of propositions found in the static context.

PL-3 Building an Ontology of Cognitive Conceptual Spaces (Demo and Poster)

Andrei Lapets, Lily Wong

Supervised By: Assaf Kfoury

In an extensive survey, Gärdenfors outlines a model for cognitive processing and semantics in which qualia, concepts, properties, and relationships are represented as regions and maps over metric spaces. We consider the problem of extracting from existing ontologies and databases found on the web, such as WordNet and ConceptNet, representations of concepts and properties that are consistent with the definitions in Gärdenfors' survey. We develop intuitive interfaces for extracting this information from the ontologies with the help of many non-expert human users. We then evaluate the quality of the information extracted in this manner by considering possible applications thereof, such as in the design of languages for user interfaces and automated reasoning assistants.

PL-4 3-approximation Online Dynamic Storage Allocation (Poster)

Likai Liu

Supervised By: Hongwei Xi

The dynamic storage allocation (DSA) problem describes how a program's runtime system allocates memory for objects. It is like bin packing, but with insertion and removal of objects over time. The chief difficulties are (1) object has various sizes, (2) allocation request cannot be divided into smaller requests, and (3) once allocated, object is immobile. The problem is NP-hard, but polynomial time approximation algorithms exist for offline DSA where allocation and deallocation sequence is known in advance. The approximation is a constant factor bound for the total memory used for allocation. The constant factor can be understood as storage space wasted due to fragmentation. Here we show that, leveraging virtual memory to provide multiple arenas for memory allocation, it is possible to construct a 3-approximation online DSA (i.e. use at most three times the memory in the worst case fragmentation).

PL-5 Specifying and verifying a sorting Algorithm in ATS (Poster)

Zhiqiang Ren

Supervised By: Hongwei Xi

Specifying and then verifying properties of modular programs properties modularly can be of great use in constructing safe and reliable software systems. ATS is a programming language equipped with a highly expressive type system that allows the programmer to specify and then verify program properties within ATS itself. In an attempt to illuminate this process, we use a sorting algorithm as an example to demonstrate how program specification and verification can be done effectively and flexibly in ATS.

Security

S-1 ACM BUILDS Project: Practical Analysis of Buffer Overflow Vulnerabilities and Privilege Escalation using gsubmit (Demo and Poster)

Kyle Brogle, David House

Reviewed By: Leo Reyzin and Peter Gacs

Buffer overflows are a commonly overlooked attack vector that allow third parties to gain unauthorized control over system resources. In this presentation we outline the methodology behind a buffer overflow attack on 'gsubmit', a popular assignment submission program used within the BU CS department. A vulnerability resulting from unsafe system calls allows us to execute arbitrary commands as root user on servers running vulnerable versions of gsubmit, resulting in a system-wide loss of security via privilege escalation. We demonstrate a proof-of-concept for this attack, and describe specific measures taken to patch this security vulnerability on newer versions of gsubmit. We also introduce "Low-hanging Fruit", a program that examines binaries and automatically identifies their susceptibility to buffer overflow attacks.

S-2 Security Analysis of Web-based Languages using Lightweight Formal Methods (Poster)

Mark Reynolds

Supervised By: Assaf Kfoury

A method is described for analyzing web-based programming languages, in particular Java and JavaScript, using lightweight formal methods. This approach represents a significant departure from traditional signature-based approaches to detecting web-based malware, which are often easy to defeat by slight mutations of the malicious code. In the current approach, a formal specification for the security properties of the underlying language is written in the modeling language Alloy. A translator then converts web code (JavaScript scripts or Java bytecodes) into Alloy as well. When the formal specification and the translated code are combined and analyzed, any counterexamples found represent security policy violations, and thus potentially malicious code. This approach has been shown to detect several different forms of real malware.

Theory

T-1 Q-Learning in Stochastic Repeated Games (Poster)

M. Ghasemi

Supervised By: S. Homer

Game theory is a well structured environment for analyzing interactive multi-agent systems and it is largely used for development of general multi-agent learning (MAL) algorithms. Most of the MAL algorithms developed so far work based on a kind of information sharing which is not practical at least in competitive games. Therefore the goal herein is to face the problem with another approach, which is updating policies in different orders of the time. This novel algorithm is shown to work well both in matrix games and stochastic infinite-games. It uses the aspects of reinforcement learning (RL) algorithms with some modifications to be useful for games. The main amendment is using different frequencies for policy updating of each agent. Using this method (QDF), the agents learn the Stackelberg equilibrium (or time-delayed Nash) strategies without having information about their rivals' actions, payoffs, and even their actions' sets. Some famous computer simulation studies are presented to show the applicability of the proposed algorithm.

T-2 Complexity of the Restricted Stable Paths Problem and Variants (Poster)

Kevin Donnelly, Andrei Lapets

Supervised By: Assaf Kfoury

Interdomain routing on the Internet is performed using route preference policies specified independently, and arbitrarily by each autonomous system (AS) in the network. These policies are used in the border gateway protocol (BGP) by each AS when selecting next-hop choices for routes to each destination. Conflicts between policies used by different ASs can lead to routing instabilities that, potentially, cannot be resolved regardless of how long BGP runs. The stable paths problem (SPP) is an abstract graph theoretic model of the problem of selecting next-hop routes for a destination. A solution to this problem is a set of next-hop choices, one for each AS, that is compatible with the policies of each AS. In a stable solution each AS has selected its best next-hop if the next-hop choices of all neighbors are fixed. BGP can be viewed as a distributed algorithm for solving an SPP instance. We present a collection of positive and negative computational complexity results for a collection of restricted variants of SPP. Several simple variants are shown to be NP-complete, and efficient (centralized) algorithms are provided for networks that satisfy natural constraints on network topology and network link weights.

T-3 Clustering Large Datasets With Limited Distance Information (Poster)

Konstantin Voevodski

Supervised By: Shang-Hua Teng, Yu Xia

We study the problem of clustering large datasets with limited distance information. Many application domains have recently faced an explosion of data, presenting several challenges to traditional clustering techniques. In particular, computing the distances between all pairs of points, as required by traditional clustering algorithms, is now often infeasible. We develop two types of solutions for this setting. The first assumes that we are given the ability to use one versus all distance queries to ask for distances between a single point and all other points in the dataset. We show that given an assumption about the utility of the k-median objective function in clustering the points, we can efficiently produce an accurate clustering using only $O(k)$ queries. The second solution uses a heuristic nearest neighbor search to build a graph from the dataset, and then performs iterative local partitioning. We provide evidence that both approaches are effective in clustering large biological datasets that contain tens of thousands of points.

T-4 Information Asymmetries in Pay-Per-Bid Auctions (Poster)

Georgios Zervas

Supervised By: John Byers, Michael Mitzenmacher

Recently, some mainstream e-commerce web sites have begun using “pay-per-bid” auctions to sell items, from video games to bars of gold. In these auctions, bidders incur a cost for placing each bid in addition to (or sometimes in lieu of) the winner’s final purchase cost. Thus even when a winner’s purchase cost is a small fraction of the item’s intrinsic value, the auctioneer can still profit handsomely from the bid fees. Our work provides novel analyses for these auctions, based on both modeling and datasets derived from auctions at Swoopo.com, the leading pay-per-bid auction site. While previous modeling work predicts profit-free equilibria, we analyze the impact of information asymmetry broadly, as well as Swoopo features such as bidpacks and the Swoop It Now option specifically. We find that even small asymmetries across players (cheaper bids, better estimates of other players’ intent, different valuations of items, committed players willing to play “chicken”) can increase the auction duration significantly and thus skew the auctioneer’s profit disproportionately. We discuss our findings in the context of a dataset of thousands of live auctions we observed on Swoopo, which enables us also to examine behavioral factors, such as the power of aggressive bidding. Ultimately, our findings show that even with fully rational players, if players overlook or are unaware any of these factors, the result is outsized profits for pay-per-bid auctioneers.