

DEPARTMENT OF COMPUTER SCIENCE
BOSTON UNIVERSITY
ABSTRACTS FOR IAP POSTER SESSION
MARCH 22, 2005

Contents

Cryptography	3
Finding Collisions on a Public Road	3
Intrusion-Resilient Secure Channels	3
Simple Privacy-Preserving Range Queries	4
Upper and Lower Bounds on Black-Box Steganography	4
Networking	5
Compositional Analysis of Networked Applications with TRAFFIC	5
A Compositional Analysis Algorithm for Flow Specification	5
Generating the Internet's router-level topology using the first-principles approach	6
Bayesian Packet Loss Detection for TCP	6
M2RC: Multiplicative-increase/additive-decrease Multipath Routing Control for Wire- less Sensor Networks	7
Boundary and Contour map Estimation in Wireless Sensor Networks	7
Exploiting the transients of adaptation for RoQ attacks on Internet Elements	7
Approximately Uniform Random Sampling in Sensor Networks	8
A Topology-Aware Scalable Network Architecture for Providing Per-Flow Quality-of- Service	9
The Effect of Router Buffer Size on HighSpeed TCP Performance	9
Mining Anomalies Using Traffic Distributions	10
Characterizing and Modeling Traffic Matrices	10
Operating and Database Systems	11
Friendly Virtual Machines: Leveraging a Feedback-Control Model for Application Adap- tation	11
End-to-end Window-Constrained Scheduling for Real-Time Communication	11
Efficient end-host multicast on the scale of the Internet	11

Network Positioning for Nearest Neighbors	12
On Trip Planning Queries in Spatial Databases	12
Discovering Frequent Arrangements of Temporal Intervals	13
Characterizing and Exploiting Reference Locality in Data Stream Applications	13
Programming Languages	14
A Formal Semantics For Weak References	14
A Typed High Level Assembly Language	14
A Functional Approach to Typed Object-Oriented Programming	15
The Efficient Generation of Random Lambda Calculus Terms	15
Safe Programming with Pointers	15
Compositional Analysis of Networked Applications with TRAFFIC	15
A Compositional Analysis Algorithm for Flow Specification	16
Theory of Computation	17
Random Strings Collapse The Polynomial Hierarchy	17
Quantum Computer: Power and Limitation	17
Cache Oblivious Algorithms for Finite Element Matrix and N-body Simulation	17
Vision and Graphics	19
Efficient and Accurate Gesture Spotting via Pruning and Subgesture Reasoning	19
Predicting where to look for motion in an active camera network	19
Non-Rigid Image Registration Using Multiple Costs	20
Human Pose Estimation Over a Smooth Embedded Space	20
Tracking Large Variable Numbers of Objects in Clutter	20
Hand Pose Estimation by a Directed Acyclic Graph of Pairwise Classifiers	20
Multi-scale 3D Scene Flow from Binocular Stereo Sequences	21
Automatic 2D Hand Tracking in Video Sequences	22
Projective invariant trajectory representation and matching	22
Learning Embeddings for Indexing and Classification	23
MosaicShape: Stochastic Region Grouping with Shape Prior	23
A Vision Interface Based on Gaze and Blink Detection	24

Cryptography

C1. Finding Collisions on a Public Road

*Presenter: Chun-Yuan Hsiao
with Leonid Reyzin*

Many cryptographic primitives begin with parameter generation, which picks a primitive from a family. Such generation can use public coins (e.g., in the discrete-logarithm-based case) or secret coins (e.g., in the factoring-based case). We study the relationship between public-coin and secret-coin collision-resistant hash function families (CRHFs). Specifically, we demonstrate that: 1) there is a lack of attention to the distinction between secret-coin and public-coin definitions in the literature, which has led to some problems in the case of CRHFs; 2) in some cases, public-coin CRHFs can be built out of secret-coin CRHFs; 3) the distinction between the two notions is meaningful, because in general secret-coin CRHFs are unlikely to imply public-coin CRHFs. The last statement above is our main result, which states that there is no black-box reduction from public-coin CRHFs to secret-coin CRHFs. Our proof for this result, while employing oracle separations, uses a novel approach, which demonstrates that there is no black-box reduction without demonstrating that there is no relativizing reduction.

C2. Intrusion-Resilient Secure Channels

*Presenter: Robert McNerney
with Scott Russell, Gene Itkis*

Suppose that Alice lives in Alaska and her friend Bob lives in Boston. In order to communicate privately, they agree upon some secret keys and set up a “secure channel” using some encryption scheme. Unfortunately, some malicious party Eve might covertly obtain the secret keys of Alice, Bob, or even both of them. From that point onward Eve will be able to decipher their supposedly private conversation. We propose a new primitive called an Intrusion-Resilient Secure Channel which automatically restores the privacy of channel messages after such a compromise by Eve. Recovery is possible even if Eve observes (passively, without disturbing) all messages exchanged by Alice and Bob. We give a formal definition of a two-party intrusion-resilient channel and provide a simple, generic construction using ordinary, existing public-key encryption schemes. We prove a concrete bound relating the security of this channel construction to that of its underlying encryption scheme. Additionally, we suggest a general strategy and conditions under which one may apply intrusion-resilient channels to obtain security improvements for more general two-party cryptographic protocols. As a case-study, we incorporate our channels into a particular intrusion-resilient digital signature scheme and show an improved security result for the augmented scheme.

C3. Simple Privacy-Preserving Range Queries

*Presenter: Scott Russell
with Nenad Dedic, Leonid Reyzin*

The ability to search public and proprietary databases via the Internet has been a boon for individual information seekers. However, in some situations an individual might like their queries and the resulting information to remain confidential even from the organization providing the database. For example, imagine a query containing an unlisted telephone number. We demonstrate a protocol for performing simple range queries of the form "What values in the table are between 100 and 800?" against a server-hosted table of numerical data in a 'privacy-preserving' manner. Privacy-preserving means if both parties follow the protocol honestly, the individual learns only the table values in the range and nothing about values outside the range. On the other hand, the data provider learns only the number of table values returned and nothing about the query range or returned values. To the best of our knowledge this is the first provably-secure privacy-preserving range query protocol. The protocol is fairly efficient taking $S + \log N$ rounds, where S is the number of values returned, and N is the total number of values in the table. The number of bits communicated in each round is proportional to $\log N$.

C4. Upper and Lower Bounds on Black-Box Steganography

*Presenter: Nenad Dedic
with Gene Itkis, Leonid Reyzin, Scott Russell*

The goal of steganography is undetectable communication. A stegosystem communicates a message in such way that all the eavesdropper observes is an innocuous-looking conversation. In order for the stegosystem to work, it must have a notion of what "innocuous-looking" means. This is formally modeled by a channel: a distribution of allowed conversations. We study so-called "black-box" stegosystems: ones that are constructed to work for any channel as long as they can sample from the distribution. Our main concern is the rate at which data can be covertly sent. We present black-box stegosystems that attain higher rates than previously known. We also prove that these rates are optimal. To that end we show that any increase in rate must result in an exponential increase in time required to encode a secret message, thereby exposing a fundamental limitation of black-box steganography.

Networking

N1. **Compositional Analysis of Networked Applications with TRAFFIC**

Presenter: Adam D. Bradley

with Azer Bestavros, Assaf Kfoury, Ibrahim Matta In spite of a wealth of theoretical tools for ensuring "good" behaviors for compositional network applications and services (e.g., the Network calculus, QoS theory, queuing theory, control theory), very few have been widely employed in practical systems or development processes, in part because of their relatively steep and unintuitive learning curves. In light of this, we present the TRAFFIC (Typed Representation and Analysis of Flows For Interoperability Checks) framework, in which compositional application components are represented as "typed flows". Flows represent any logical component of a network application, from program control to schedulers to transport networks. The points at which flows interact (sockets) are assigned types which represent simplified statements from theories of compositional analysis; the TRAFFIC type inference system is then able to assess whether particular flows can be composed without violating desirable invariants of performance or correctness.

N2. **A Compositional Analysis Algorithm for Flow Specification**

Presenter: Likai Liu

with Azer Bestavros, Assaf Kfoury, Ibrahim Matta

Compositional analysis is a type inference algorithm that infers principal typing of a program—the most general typing of which all other typings are an instance, subject to substitution. This is useful for separate compilation where modular and incremental analysis is a requirement. We devise a compositional analysis algorithm for a specific kind of program, written in the language TRAFFIC (see joint work with Adam Bradley), that represents "flows" of a network of controllers, where type checking of such a program involves the use of network flow control theory.

N3. Generating the Internet's router-level topology using the first-principles approach

*Presenter: Chong Wang
with John Byers*

Considerable recent research has focused on developing a better understanding the Internet's router-level topological structure: both on developing better maps of the topology and on devising more accurate models of router-level network growth. While most proposed models focus exclusively on statistical properties, a recently proposed "first-principles" approach developed at Caltech and AT&T considers a theory of router-level topology growth that reflects practical technological and economic constraints and tradeoffs. Following this first-principles approach, we explore the details of optimization processes in the construction of networks and generation of network topologies. We formulate the network design problem in an optimization framework and design heuristics to yield approximately optimal outcomes. Our ongoing work includes validation of this topology generation method with real router-level network topologies.

N4. Bayesian Packet Loss Detection for TCP

*Presenter: Nahur Fonseca
with Mark Crovella*

One of TCP's critical tasks is to determine which packets are lost in the network, as a basis for control actions (flow control and packet retransmission). Modern TCP implementations use two mechanisms: timeout, and fast retransmit. Detection via timeout is necessarily a time-consuming operation; fast retransmit, while much quicker, is only effective for a small fraction of packet losses. In this paper we consider the problem of packet loss detection in TCP more generally. We concentrate on the fact that TCP's control actions are necessarily triggered by inference of packet loss, rather than conclusive knowledge. This suggests that one might analyze TCP's packet loss detection in a standard inferencing framework based on probability of detection and probability of false alarm. This paper makes two contributions to that end: First, we study an example of more general packet loss inference, namely optimal Bayesian packet loss detection based on round trip time. We show that for long-lived flows, it is frequently possible to achieve high detection probability and low false alarm probability based on measured round trip time. Second, we construct an analytic performance model that incorporates general packet loss inference into TCP. We show that for realistic detection and false alarm probabilities (as are achievable via our Bayesian detector) and for moderate packet loss rates, the use of more general packet loss inference in TCP can improve throughput by as much as 25%.

N5. M2RC: Multiplicative-increase/additive-decrease Multipath Routing Control for Wireless Sensor Networks

Presenter: Hanny Morcos

with Azer Bestavros and Ibrahim Matta

Routing protocols in wireless sensor networks (WSN) face two main challenges: first, the challenging environments in which WSN's are deployed negatively affect the quality of the routing process. Therefore, routing protocols for WSN's should recognize and react to node failures and packet losses. Second, sensor nodes are battery-powered, which makes power a scarce resource. Routing protocols should optimize power consumption to prolong the lifetime of the WSN. In this paper, we present a new adaptive routing protocol for WSN's, we call it M2RC. M2RC has two phases: mesh establishment phase and data forwarding phase. In the first phase, M2RC establishes the routing state to enable multipath data forwarding. In the second phase, M2RC forwards data packets from the source to the sink. Targeting hop-by-hop reliability, an M2RC forwarding node waits for an acknowledgment (ACK) that its packets were correctly received at the next neighbor. Based on this feedback, an M2RC node applies multiplicative-increase/additive-decrease (MIAD) to control the number of neighbors targeted by its packet broadcast. We simulated M2RC in the ns-2 simulator [4] and compared it to GRAB [1], Max-power, and Min-power routing schemes. Our simulations show that M2RC achieves the highest throughput with at least 10-30 percent less consumed power per delivered report in scenarios where a certain number of nodes unexpectedly fail.

N6. Boundary and Contour map Estimation in Wireless Sensor Networks

Presenter: Niky Riga

with Ibrahim Matta and Azer Bestavros Wireless sensor networks have emerged as an enabling tool for monitoring inaccessible environments; surveillance in military zones, habitat monitoring, evaluation of building structures etc. These networks are envisioned to consist of many small devices each with limited capabilities. Although each node has power, computation and communication constraints the network as a whole is envisioned to support a wide range of demanding applications. An important set of queries for sensor networking applications is those concerning boundary estimation. We refer to boundaries as a set of imaginary curves that split the field into regions. All the sensors within a region share some common characteristic, e.g. detect the same event, have similar sensing values. The problem of boundary estimation is determining which sensors lie within a distance r from these imaginary lines. Recent work addressed this problem through statistical analysis of neighboring information gathered through message exchange. A drawback for this approach is high communication cost. Other approaches are based on hierarchical structures on the sensor field, which are less robust and scalable. We are currently investigating the tradeoffs between existing techniques and proposing an efficient, distributed and probabilistic protocol.

N7. Exploiting the transients of adaptation for RoQ attacks on Internet Elements

Presenter: Mina Guirguis

with Azer Bestavros and Ibrahim Matta Current computing systems depend on adaptation mechanisms to drive the system into quiescent regions of operation. These regions are often characterized by being efficient, fair and stable. Research studies have focused on developing more sophisticated adaptation mechanisms without the proper attention to their security aspects. In this work, we expose vulnerabilities in adaptation mechanisms against Reduction of Quality (RoQ) attacks. RoQ attacks are a special form of attacks that target the transients of a system's adaptive behavior as opposed to its limited steady-state capacity. We show that a well orchestrated attack could introduce significant inefficiencies that could potentially deprive an Internet element from much of its capacity, or significantly reduce its service quality, while evading detection by consuming an unsuspecting, small fraction of that system's hijacked capacity. We develop a control theoretic model for assessing the impact of RoQ attacks on a number of systems, including load balancers, admission controllers and Internet congestion control protocols. We quantify the damage inflicted by an attacker through deriving appropriate metrics in each setting. We validate our findings through real Internet experiments performed in our lab.

N8. Approximately Uniform Random Sampling in Sensor Networks

Presenter: Boulat A. Bash

with Jef Considine, John Byers

Recent work in sensor databases has focused extensively on distributed query problems, notably distributed computation of aggregates such as average, count, and sum. Existing methods for computing aggregates broadcast queries to all sensors and use in-network aggregation of responses to minimize messaging costs. In this work, we focus on distributed algorithms to sample each sensor node in the network with equal probability. Such uniform random sampling across nodes can serve both as an alternative building block for aggregation and as an integral component of many other useful randomized algorithms. Prior to our work, the best existing proposals for uniform random sampling of sensors involve contacting all nodes in the network. We propose a practical method which is only approximately uniform, but contacts a number of sensors proportional to the diameter of the network instead of its size. The approximation achieved is tunably close to exact uniform sampling, and only relies on well-known existing primitives, namely geographic routing, distributed computation of Voronoi regions and von Neumann's rejection method. Ultimately, our sampling algorithm has the same worst-case asymptotic cost as routing a point-to-point message, and thus it is asymptotically optimal among request/reply-based sampling methods. We provide experimental results demonstrating the effectiveness of our algorithm on both synthetic and real sensor topologies.

N9. A Topology-Aware Scalable Network Architecture for Providing Per-Flow Quality-of-Service

Presenter: Kanishka Gupta

with John Byers and Azer Bestavros

Providing scalable fair-queuing solutions has been a well-studied research problem in computer networks in the past decade. Traditional architectures have provided either scalability (FIFO), by being stateless, or rich services (Intserv), by providing per-flow state, but not both. Recently proposed architectures such as State-less core (SCORE) have overcome this problem by making the packets, rather than the routers, carry the state. The routers in the core network then use efficient distributed algorithms which utilize this information to provide per-flow service. Many improvements to SCORE have been proposed which essentially build up upon the idea of marking packets at the edges. In this work, we investigate a completely different architecture which also achieves the same result of providing per-flow fair share without keeping per-flow state. Our approach is somewhat similar in spirit to the Diff-Serv architecture, in the sense that we maintain per-class state at the core-routers. The key idea, however, is that unlike Diffserv, we classify the flows into the classes dynamically based on the network-topology, rather than statically based on the flow's service class. Being cognizant of the topology enables us to provide per-flow share with only a small (constant) number of classes equal to the in-degree of a router. Our definition of class makes the network into a hierarchy and in this respect is similar to previous work done on hierarchical packet scheduling.

N10. The Effect of Router Buffer Size on HighSpeed TCP Performance

Presenter: Georgios Smaragdakis

with Dhiman Barman, Mark Crovella

We study the effect of the IP router buffer size on the throughput of HighSpeed TCP (HSTCP). We are motivated by the fact that in high speed routers, the buffer size is important as such a large buffer size might be a constraint. We first derive an analytical model for HighSpeed TCP and we show that for small buffer size equal to 10% of the bandwidth-delay product, HighSpeed TCP can achieve more than 90% of the bottleneck capacity. We also show that setting the buffer size equal to 20% can increase the utilization of HighSpeed TCP up to 98%. On the contrary, setting the buffer size to less than 10 percent of the bandwidth-delay product can decrease HighSpeed TCP's throughput significantly. We also study the performance effects under both DropTail and RED AQM. Analytical results obtained using a fixed-point approach are compared to those obtained by simulation.

N11 Mining Anomalies Using Traffic Distributions

*Presenter: Anukool Lakhina
with Mark Crovella and Christophe Diot*

The increasing practicality of large-scale flow capture makes it possible to conceive of traffic analysis methods that detect and identify a large and diverse set of anomalies. However the challenge of effectively analyzing this massive data source for anomaly diagnosis is as yet unmet. We argue that the distributions of packet features (IP addresses and ports) observed in flow traces reveals both the presence and the structure of a wide range of anomalies. Using entropy as a summarization tool, we show that the analysis of feature distributions leads to significant advances on two fronts: (1) it enables highly sensitive detection of a wide range of anomalies, augmenting detections by volume-based methods, and (2) it enables automatic classification of anomalies via unsupervised learning. We show that using feature distributions, anomalies naturally fall into distinct and meaningful clusters. These clusters can be used to automatically classify anomalies and to uncover new anomaly types. We validate our claims on data from two backbone networks (Abilene and Geant) and conclude that feature distributions show promise as a key element of a fairly general network anomaly diagnosis framework.

N12 Characterizing and Modeling Traffic Matrices

*Presenter: Vijay Erramili
with Mark Crovella*

There exist a wide variety of network design problems that require a traffic matrix as input to carry out performed evaluation. For instance, in order to assess the performance of routing protocols, load balancers, failure recovery mechanisms, one typically needs to have some information on the traffic demands in the network. However due to a variety of reasons, there haven't been any realistic models of traffic matrices or any way to synthesize traffic matrices without considering often unrealistic assumptions like spatial independence of source/destination nodes. In this study we present a first attempt to model traffic matrices by first characterizing them and by proposing simple models. We have observed a property of "topological locality" in traffic that suggests that nearby nodes are more likely to exchange traffic than would be predicted by a purely independent model. While this property has been noted since the earliest days of the ARPANET, it has not been previously characterized quantitatively. We show that deviations from the pure independence model, including deviations due to "locality", must obey certain constraints which allows us to propose models for topological locality, and how this can help us in constructing useful models for traffic matrices. We consider multiple data sets of traffic matrices over three different networks, two of them educational (Abilene and Geant) and one commercial (Sprint) for our investigations.

Operating and Database Systems

01. **Friendly Virtual Machines: Leveraging a Feedback-Control Model for Application Adaptation**

Presenter: Yuting Zhang

with Mina Guirguis, Azer Bestavros, Ibrahim Matta and Richard West

With the increased use of “Virtual Machines” (VMs) as vehicles that isolate applications running on the same host, it is necessary to devise techniques that enable multiple VMs to share underlying resources both fairly and efficiently. To that end, one common approach is to deploy complex resource management techniques in the hosting infrastructure. Alternately, in this paper, we advocate the use of self-adaptation in the VMs themselves based on feedback about resource usage and availability. Consequently, we define a “Friendly” VM (FVM) to be a virtual machine that adjusts its demand for system resources, so that they are both efficiently and fairly allocated to competing FVMs. Such properties are ensured using one of many provably convergent control rules, such as AIMD. By adopting this distributed application-based approach to resource management, it is not necessary to make assumptions about the underlying resources nor about the requirements of FVMs competing for these resources. To demonstrate the elegance and simplicity of our approach, we present a prototype implementation of our FVM framework in User-Mode Linux (UML)—an implementation that consists of less than 500 lines of code changes to UML. We present an analytic, control-theoretic model of FVM adaptation, which establishes convergence and fairness properties. These properties are also backed up with experimental results using our prototype FVM implementation.

02. **End-to-end Window-Constrained Scheduling for Real-Time Communication**

Presenter: Yuting Zhang

with Richard West

This paper extends our original work on window-constrained scheduling, to address the problem of meeting end-to-end service guarantees across a sequence of servers. We describe an algorithm, called Multi-hop Virtual Deadline Scheduling (MVDS), that attempts to minimize end-to-end window-constraint violations, while maximizing link utilization for a series of real-time streams. The challenge posed by the multi-hop problem is how to derive a local scheduling and dropping scheme from global service requirements, so that each server along a path can cooperate to guarantee end-to-end service. Similar to our VDS algorithm developed for a single server, MVDS orders packets at the heads of streams according to their local virtual deadlines. Using various packet dropping schemes at each server, based on current workloads and likelihoods of meeting end-to-end service constraints, we evaluate the performance of MVDS. Simulation results show that MVDS can provide better window-constrained service guarantees than other related algorithms, while still maintaining high link utilization.

03. Efficient end-host multicast on the scale of the Internet

*Presenter: Gabriel Parmer
with Richard West*

Applications that stream data over the Internet to a large group of subscribers require a means to create multicast trees over which the data will be transported. Given the relatively small level of IP-level multicast adoption, end-host multicast tree construction is a viable method. Much research has been carried out regarding the use of overlay topologies to create these trees because overlays provide reliability and scalability. We propose a novel method for the construction of end-host multicast trees on top of an overlay which is scalable for the common operations, and distributes links stress, yet still provides good latency of delivery to the subscribers.

04. Network Positioning for Nearest Neighbors

*Presenter: Xin Qi
with Dihan Cheng, Richard West and Shang-hua Teng*

Current work on the network positioning problem focuses on geometric approaches to preserve real distances in a corresponding embedded space. However, rather than determining actual distances between hosts, it is usually only necessary to determine nearest neighbors. Applications such as peer-to-peer systems and content-delivery networks can benefit from positioning techniques that estimate nearest neighbors without direct measurement of all physical distances between hosts. In this paper, we find that among different normed vector spaces, $\{L_p \mid p \in [1, \infty)\}$, there exists a linear relationship with respect to the effect of an on-line refinement algorithm, such that L_∞ is the most accurate for preserving real distance. This enables hosts to be iteratively eliminated when probing for nearest neighbors. Although the triangle inequality must hold in embedding theory, there are cases where it fails for real networks. We provide a simple approach to compensate for violations of this inequality when estimating distances. We show that with the aid of an off-line network positioning mechanism based on the L_∞ norm, using an on-line refinement algorithm to adjust the prediction, a host in the network can effectively find its nearest neighbor or neighbor set among a large number of candidates.

05. On Trip Planning Queries in Spatial Databases

*Presenter: Feifei Li
with Dihan Cheng, Marios Hadjieleftheriou, George Kollios, and Shang-Hua Teng* In this paper we discuss a new type of query in Spatial Databases, called the Trip Planning Query (TPQ). Given a set of points P in space, where each point belongs to a specific category, and two points S and E , TPQ retrieves the best trip that starts at S , passes through at least one point from each category, and ends at E . The problem is NP-hard, since it is a generalization of the Traveling Salesman Problem. The difficulty of this query lies in the existence of multiple choices per category. In this paper, we study fast approximation algorithms for TPQ

in a metric space. We provide a number of approximation algorithms with approximation ratios that depend on either the number of categories, the maximum number of points per category or both. Therefore, for different instances of the problem, we can choose the algorithm with the best approximation ratio, since they all run in polynomial time. Furthermore, we use some of the proposed algorithms to derive efficient heuristics for large datasets stored in external memory. Finally, we give an experimental evaluation of the proposed algorithms using synthetic datasets generated on real road networks.

06. **Discovering Frequent Arrangements of Temporal Intervals**

Presenter: Panagiotis Papapetrou

with George Kollios, Dimitrios Gunopoulos and Stan Sclaroff

A new problem in temporal pattern mining is studied: how to find frequent arrangements of temporal intervals. It is assumed that the database consists of sequences of events, where an event occurs during a time-interval. The goal is to mine general temporal arrangements of event intervals that appear frequently in the database. Since in practice most events are not instantaneous but occur over a period of time, and since different events may occur concurrently, there are many practical applications that require mining such temporal correlations between intervals. Efficient methods to find frequent arrangements of temporal intervals using both breadth first and depth first search techniques are described. The performance of the proposed algorithms is evaluated and compared with other approaches on real datasets (American Sign Language streams) and large synthetic datasets.

07. **Characterizing and Exploiting Reference Locality in Data Stream Applications**

Presenter: Ching Chang

with Feifei Li, George Kollios, Azer Bestavros

Previous research literature on data stream management has been limited to describing systems which depend upon conventional IID assumption. Our study identifies that IID alone fails to characterize many real data sets. We then propose a novel approach to conceptualize the dynamics in data stream by describing two causes of reference locality: popularity over long time scales and temporal correlations over shorter time scales. An elegant mathematical model is adopted to precisely quantify the degree of those sources of locality. We then analyze the impact of locality-awareness on achievable performance gains over traditional algorithms on applications such as MAX-subset join and Lossy Counting. Finally, the concept of entropy is introduced in the context of data stream as a measurement for space usage in summarization applications. Our experiments using both real and synthetic data sets validate our ideas and analysis.

Programming Languages

P1. A Formal Semantics For Weak References

*Presenter: Joe Hallett
with Assaf Kfoury*

A weak reference is a reference to an object that is not followed by the pointer tracer when garbage collection is called. That is, a weak reference cannot prevent the object it references from being garbage collected. Weak references remain a troublesome programming feature largely because there is not an accepted, precise semantics that describes their behavior (in fact, we are not aware of any formalization of their semantics). The trouble is that weak references allow reachable objects to be garbage collected, therefore allowing garbage collection to influence the result of a program. Despite this difficulty, weak references continue to be used in practice for reasons related to efficient storage management, and are included in many popular programming languages (Standard ML, Haskell, OCaml, and Java). We give a formal semantics for a calculus called lambda-weak that includes weak references and is derived from Morrisett, Felleisen, and Harper's lambda-gc. lambda-gc formalizes the notion of garbage collection by means of a rewrite rule. Such a formalization is required to precisely characterize the semantics of weak references. However, the inclusion of a garbage-collection rewrite-rule in a language with weak references introduces non-deterministic evaluation, even if the parameter-passing mechanism is deterministic (call-by-value in our case). This raises the question of confluence for our rewrite system. We discuss methods in which we can recover deterministic evaluation of programs. We define conditions that allow other garbage collection algorithms to co-exist with our semantics of weak references. We also introduce a polymorphic type system to prove the absence of erroneous program behavior (i.e., the absence of "stuck evaluation") and a corresponding type inference algorithm. We prove the type system sound and the inference algorithm sound and complete.

P2. A Typed High Level Assembly Language

*Presenter: Sa Cui
with Hongwei Xi*

ATS is a recently developed programming language with an advanced type system rooted in the framework Applied Type System (ATS). In ATS, a variety of programming paradigms are supported in a typeful manner, including functional programming, object-oriented programming and imperative programming with pointers. We present a typed high level assembly language (THLA) in the framework ATS, the type system of which supports a novel notion of stateful views which is used to describe the state of registers, heap and stack. This work makes it possible to support a style of assembly programming in ATS.

P3. A Functional Approach to Typed Object-Oriented Programming

Presenter: Rui Shi

with Chiyen Chen, Hongwei Xi

ATS is a recently developed programming language with a highly expressive type system rooted in the framework Applied Type System. In this paper, we present a design to support typed object-oriented programming (OOP) with multiple inheritance in ATS. In contrast to most existing approaches that represent objects as records, we instead represent objects as functions in ATS. We demonstrate that full-fledged support for OOP can be built directly on the top of the functional core of ATS without seeking ad hoc extensions of ATS. In particular, we show that a variety of issues with OOP (e.g., binary methods, the self type, parametric polymorphism, multiple inheritance) can be readily addressed by the approach we take to support OOP.

P4. The Efficient Generation of Random Lambda Calculus Terms

Presenter: Jue Wang

with Assaf Kfoury, Franklyn Turbak (Wellesley College) We explore the problem of generating lambda calculus terms of a given size uniformly at random. This work has several motivations. First, through performing statistical sampling experiments with randomly generated terms, we can study interesting properties of lambda calculus terms. Second, random lambda calculus terms can serve as inputs to program analysis algorithms such as type inference to evaluate both their performance and correctness. We present an algorithm that generates a random lambda calculus term of a given size, assuming uniform distribution over all terms of a given size. To improve the efficiency of generating a term, the current algorithm makes use of memoization techniques and also employs a system of number representation that represents numbers approximately. In addition, we also present some of the possible applications for such a tool with some preliminary results.

P5. Safe Programming with Pointers

Presenter: Dengping Zhu

with Hongwei Xi

We present a type system that can effectively facilitate the use of types in capturing invariants in stateful programs that may involve (sophisticated) pointer manipulation. With its root in a recently developed framework Applied Type System (ATS), the type system imposes a level of abstraction on program states by introducing a novel notion of recursive stateful views and then relies on a form of linear logic to reason about such views. We consider the design and then the formalization of the type system to constitute the primary contribution of the paper. In addition, we mention a prototype implementation of the type system and then give a variety of examples that attests to the practicality of programming with recursive stateful views.

P6/N1. Compositional Analysis of Networked Applications with TRAFFIC

Presenter: Adam D. Bradley

with Azer Bestavros, Assaf Kfoury, Ibrahim Matta

In spite of a wealth of theoretical tools for ensuring "good" behaviors for compositional network applications and services (e.g., the Network calculus, QoS theory, queuing theory, control theory), very few have been widely employed in practical systems or development processes, in part because of their relatively steep and unintuitive learning curves. In light of this, we present the TRAFFIC (Typed Representation and Analysis of Flows For Interoperability Checks) framework, in which compositional application components are represented as "typed flows". Flows represent any logical component of a network application, from program control to schedulers to transport networks. The points at which flows interact (sockets) are assigned types which represent simplified statements from theories of compositional analysis; the TRAFFIC type inference system is then able to assess whether particular flows can be composed without violating desirable invariants of performance or correctness.

P7/N2. A Compositional Analysis Algorithm for Flow Specification

Presenter: Likai Liu

with Azer Bestavros, Assaf Kfoury, Ibrahim Matta

Compositional analysis is a type inference algorithm that infers principal typing of a program—the most general typing of which all other typings are an instance, subject to substitution. This is useful for separate compilation where modular and incremental analysis is a requirement. We devise a compositional analysis algorithm for a specific kind of program, written in the language TRAFFIC (see joint work with Adam Bradley), that represents "flows" of a network of controllers, where type checking of such a program involves the use of network flow control theory.

Theory of Computation

T1. **Random Strings Collapse The Polynomial Hierarchy**

*Presenter: Benjamin Hescott
with Steven Homer*

Here we show that the set of resource bounded random strings create a relativized world where the polynomial hierarchy is finite when the resource is greater than polynomial space. We show that this is an easy consequence of Karp and Lipton's proof that it is unlikely that NP has small circuits. We then collapse the hierarchy to a lower level using the derandomization techniques of Allender et. al. by showing that the set of resource bounded (space) random strings are complete under nondeterministic reductions for the classes where they reside - making them fixed points for the NP jump. We also investigate how resource bounded measure and generalized betting strategies may show that a collapse to the bottom level is impossible. Thus creating a relativized world where the polynomial hierarchy is finite and P is not equal to NP.

T2. **Quantum Computer: Power and Limitation**

*Presenter: Debajyoti Bera
with Steven Homer*

The Quantum Computing model was formulated by Richard Feynmann to capture the computational aspects involved in quantum mechanics. After David Deutsch first showed how the Quantum Computing model can be used to reduce computation time for a combinatorial problem, Quantum Computing became an active area of research in theoretical Computer Science. One of the main initial focuses was solving NP-complete problems. However, to date, there are only a few results which give a quantum polynomial time algorithm for a problem which has no known classical polynomial time algorithm - and none of them involve solving an NP-complete problem. Here we investigate the power and limitations of Quantum Computing, as known till date. We show why Quantum Computing does not solve NP-complete problems by straight forward brute force searching. Yet Quantum Computing indeed helps in speeding up well-studied polynomial time algorithms. We discuss a few of them and motivate the use of quantum-aware data structures for better quantum algorithms.

T3. **Cache Oblivious Algorithms for Finite Element Matrix and N-body Simulation**

*Presenter: Kebin Wang
with Shanghua Teng*

In this paper, we present optimal cache-aware and cache-oblivious algorithms for matrix-vector multiplication of sparse matrices derived from well-shaped meshes. The key problem is to lay out the mesh to minimize cache misses. If n is the graph size, B is the cache block size, and M is the size of cache, and d is the dimensionality of the geometric domain of the mesh then an optimal layout allows the matrix-vector multiplication to run in $O(n/B)$

memory transfers in the cache-oblivious model. This result requires the "tall cache assumption" $M = \Omega(B^d)$. Our cache-aware results are as follows: A simple divide-and-conquer layout is only within a factor of $B^{1/d}$ of optimal in the DAM model, but becomes optimal with the "taller cache" assumption $M = \Omega(B^{d+1})$. Using the multi-way graph partitioning of Kiwi *et al*, an optimal cache-aware layout can be constructed with $O((n/B) \log(n/M))$ memory transfers in the cache-oblivious model. To our knowledge, this is the first optimal cache-aware solution to the Graph-Neighborhood problem. Our cache-oblivious results are as follows: Using fully-balanced decomposition trees, originally developed for VLSI theory, an optimal layout can be constructed for the cache-oblivious model, $O((n/B) \log(n/M))$ memory transfers in the cache-oblivious model. These results provide optimal cache-aware and cache-oblivious algorithms for the Conjugate Gradient method for finite-element computation, and provide good cache-oblivious algorithms for N-body simulation.

Vision and Graphics

V13. **Efficient and Accurate Gesture Spotting via Pruning and Subgesture Reasoning**

Presenter: Jonathan Alon

with Vassilis Athitsos

Vision-based recognition of gestures in continuous video streams can facilitate more natural human-computer interaction. Gesture spotting is the challenging task of locating the start and end frames of the video stream that correspond to a gesture of interest, while at the same time rejecting non-gesture motion patterns. We present a new gesture spotting and recognition algorithm that is based on the widely used continuous dynamic programming (CDP) algorithm. Our first contribution is a pruning method that allows the system to evaluate a relatively small number of hypotheses compared to CDP. Pruning is implemented by a set of model-dependent classifiers, that are learned from training examples. In our experiments, the proposed CDP with pruning was an order of magnitude faster compared to the original CDP algorithm, and recognition accuracy improved by 7%. The second contribution of the proposed gesture spotting algorithm is a subgesture reasoning process that models the fact that some gesture models can falsely match parts of other longer gestures. In our experiments, using the proposed subgesture modeling improved recognition accuracy by an additional 12%.

V2. **Predicting where to look for motion in an active camera network**

Presenter: Ugur Murat Erdem

with Stan Sclaroff

A framework is proposed that answers the following question: if a moving object is observed by one camera in a pan-tilt-zoom (PTZ) camera network, what other camera(s) might be foveated on that object within a predefined time window, and what would be the corresponding PTZ parameter settings? No calibration is assumed, and there are no restrictions on camera placement or initial parameter settings. The framework accrues a predictive model over time. To start out, the cameras follow randomized "tours" in discretized PTZ space. If a moving object is detected in the field of view of more than one camera at a particular instant or with a predefined time window, then the model is updated to record the cameras' associations and the corresponding parameter settings. As more and more moving objects are observed, the model adapts and the most frequent associations are discovered. The formulation also allows for verification of its predictions, and reinforces its correct predictions. The formulation is demonstrated in tracking people in an office environment with a three PTZ camera network.

V3. Non-Rigid Image Registration Using Multiple Costs

*Presenter: William Mullaly
with Margrit Betke*

Locally controlled non-rigid image registration methods are currently in high demand to help scientists model and understand the complex behaviors observed in their image data. For example, radiologists, medical physicist, and biomedical scientists are all interested in modeling the non-rigid deformations of the lung as seen in chest computed tomography (CT) scans. Current registration techniques rely on explicitly segmented features or the minimization of a single cost function. However, the similarity function that optimally registers ribs in a chest CT may not be the same as the best similarity function for relating healthy lung tissue to diseased lung tissue. We present our initial investigation into incorporating multiple cost functions into a single registration problem and show results on chest CT scans.

V4. Human Pose Estimation Over a Smooth Embedded Space

*Presenter: Tai-Peng Tian
with Stan Sclaroff*

Human pose estimation is an important precursor to human motion tracking and human activity recognition in video sequences. This work proposes a new framework for the human pose estimation problem. Leveraging upon advances made in the machine learning community, a new probabilistic embedding technique is used to compactly represent the prior knowledge on human poses. Human pose estimation is performed by searching over the smooth embedded space using standard optimization techniques. Experiments were performed to compare the performance between the new algorithm and the existing Specialized Mapping Architecture (SMA). Results show that the average error is at least an order of magnitude less than SMA.

V5. Tracking Large Variable Numbers of Objects in Clutter

*Presenter: Angshuman Bagchi
with Margrit Betke*

We propose a statistical data association method for visual tracking of enormously large numbers of objects. We do not assume any prior knowledge about the numbers involved, and the objects may appear or disappear anywhere in the image frame and at any time in the sequence. Our method combines the techniques of multitarget track initiation, recursive Bayesian tracking, clutter modeling, and multiple hypothesis filtering. The original multiple hypothesis filter addresses an NP-hard problem and is thus not practical. We propose pruning approaches that are linear in the number of tracked objects. We validated the effectiveness and scalability of the proposed method in experiments with video data of wildlife. In one experiment, 7,056 animals were tracked through 9,139 frames with a 99% agreement with the manually established ground truth of 7,007 animals.

V6. Hand Pose Estimation by a Directed Acyclic Graph of Pairwise Classifiers

*Presenter: Ashwin Thangali
with Quan Yuan, Stan Sclaroff*

Obtaining both efficient and accurate hand pose classification is particularly challenging, since the number of pose classes is quite large in general. It seems unavoidable that there should be a trade off between computational complexity and accuracy to gain a practical system. In this paper, a new method is proposed that recognizes hand poses by combining pairwise classifiers into a Directed Acyclic Graph (DAG). It is shown that in theory, the DAG gives comparable classification accuracy when compared with a group of traditional schemes under similar pairwise classification error; however, the DAG is more efficient. Moreover, it is shown empirically that the DAG formulation yields better classification accuracy when compared with a traditional one-versus-all scheme. A set of polar frequency features is proposed that are orientation-invariant while still providing discriminative information for a diverse range of hand poses. These frequency features are used in learning each pairwise classifier via AdaBoost. Training data sets are generated using realistic computer graphics models of human hands, with perturbations of illumination, viewing angle, pose parameters, and resolution. The resulting system is user-independent. In experiments on hundreds of real hand images, the DAG hand pose classifier gives accuracy that is better than one-versus-all, and comparable to a traditional voting scheme constructed from the same strong pairwise classifiers.

V7. Multi-scale 3D Scene Flow from Binocular Stereo Sequences

*Presenter: Rui Li
with Stan Sclaroff*

Scene flow methods estimate the three-dimensional motion field for points in the world, using multi-camera video data. Such methods combine multi-view reconstruction with motion estimation approaches. This paper describes an alternative formulation for dense scene flow estimation that provides convincing results using only two cameras by fusing stereo and optical flow estimation into a single coherent framework. To handle the aperture problems inherent in the estimation task, a multi-scale method along with a novel adaptive smoothing technique is used to gain a regularized solution. This combined approach both preserves discontinuities and prevents over-regularization - two problems commonly associated with basic multi-scale approaches. Internally, the framework generates probability distributions for optical flow and disparity. Taking into account the uncertainty in the intermediate stages allows for more reliable estimation of the 3D scene flow than standard stereo and optical flow methods allow. Experiments with synthetic and real test data demonstrate the effectiveness of the approach.

V8. Automatic 2D Hand Tracking in Video Sequences

Presenter: Quan Yuan

with Vassilis Athitsos, Stan Sclaroff

In gesture and sign language video sequences, hand motion tends to be rapid, and hands frequently appear in front of each other or in front of the face. Thus, hand location is often ambiguous, and naive color-based hand tracking is insufficient. To improve tracking accuracy, some methods employ a prediction-update framework, but such methods require careful initialization of model parameters, and tend to drift and lose track in extended sequences. In this paper, a temporal filtering framework for hand tracking is proposed that can initialize and reset itself without human intervention. In each frame, simple features like color and motion residue are exploited to identify multiple candidate hand locations. The temporal filter then uses the Viterbi algorithm to select among the candidates from frame to frame. The resulting tracking system can automatically identify video trajectories of unambiguous hand motion, and detect frames where tracking becomes ambiguous because of occlusions or overlaps. Experiments on video sequences of several hundred frames in duration demonstrate the systems ability to track hands robustly, to detect and handle tracking ambiguities, and to extract the trajectories of unambiguous hand motion.

V9. Projective invariant trajectory representation and matching

Presenter: Walter Nunziati

with Quan Yuan, Stan Sclaroff

The method presented in this paper aims to match trajectories of moving objects across uncalibrated video streams. It is assumed that trajectories lie on surfaces that can be locally approximated with a plane, and that the data are discrete time ordered sequences of image coordinates. A method based on projective invariant features measured at each observed point of the trajectory is presented. Raw data are first locally approximated with a cubic spline via least squares fitting. For each sampled point of the obtained curve, the feature is computed using a small number of points in its neighborhood, and the resulting sequence of invariant features computed along the entire trajectory forms the view invariant descriptor of the trajectory itself. Features are based on the cross ratio of five coplanar points, and time parametrization has been exploited to compute them without ambiguity due to point ordering. Similarity between descriptors of different trajectories is measured with a distance that takes into account the statistical properties of the cross ratio, and its symmetry with respect to the point at infinity. In experiments, an overall correct classification rate of about 95% has been obtained on a dataset of 58 trajectories of players in a soccer game video sequence, and an overall correct classification rate of about 80% has been obtained on matching partial segments of trajectories collected from two views with significant overlap of outdoor scenes with moving people and cars.

V10. Learning Embeddings for Indexing and Classification

*Presenter: Vassilis Athitsos
with Stan Sclaroff*

Large databases of patterns have emerged as an important tool for storing information in a wide variety of domains, from computer vision to speech recognition and bioinformatics. Such large databases can capture the wide variability of the data that we observe in the real world, and allow us to reliably recognize new patterns using nearest neighbor classification. At the same time, finding the nearest neighbors can often be too slow to be practical, especially in domains that use computationally expensive similarity measures. Examples of such measures include dynamic time warping for time series, shape context matching for edge images, or the edit distance for strings and biological sequences. BoostMap is an embedding algorithm that can significantly speed up retrieval in such spaces. Patterns are embedded into a real vector space, where distances can be measured rapidly using the Manhattan distance. The mapping is explicitly optimized to preserve a large amount of nearest neighbor information from the original space. The key novelty of BoostMap is that it formulates embedding construction as a machine learning task. This formulation allows the use of powerful machine learning methods (namely AdaBoost) for embedding optimization. In experiments with different datasets, including time series and images of handwritten digits, BoostMap significantly speeds up nearest neighbor retrieval and classification, and compares favorably to existing embedding methods.

V11. MosaicShape: Stochastic Region Grouping with Shape Prior

*Presenter: Jingbin Wang
with Margrit Betke*

A novel method that combines shape-based object recognition and image segmentation is proposed for shape retrieval from images. Given a shape prior represented in a multi-scale curvature form, the proposed method identifies the target objects in images by grouping oversegmented image regions. The problem is formulated in a unified probabilistic framework and solved by a stochastic Markov Chain Monte Carlo (MCMC) mechanism. By this means, object segmentation and recognition are accomplished simultaneously. Within each sampling move during the simulation process, probabilistic region grouping operations are influenced by both the image information and the shape similarity constraint. The latter constraint is measured by a partial shape matching process. A generalized parallel algorithm, combined with a large sampling jump and other implementation improvements, greatly speeds up the overall stochastic process. The proposed method supports the segmentation and recognition of multiple occluded objects in images. Experimental results are provided for both synthetic and real images.

V12. A Vision Interface Based on Gaze and Blink Detection

*Presenter: John Magee
with Margrit Betke*

In cases of paralysis so severe that a person's ability to control movement is limited to the muscles around the eyes, eye movements or blinks are the only way for the person to communicate. Interfaces that assist in such communication are often intrusive, require special hardware, or rely on active infrared illumination. A non-intrusive communication video-based interface system called EyeKeys was therefore developed that runs on a consumer-grade computer with video input from an inexpensive USB camera and works without special lighting. The system detects and tracks the person's face using multi-scale template correlation. Symmetry between left and right eyes is exploited to detect if the person is looking at the camera, or to the left or right side. Blinks are detected with by comparing sample templates of open and closed eyes with the eyes detected in the video images. When the closed eye template correlates better than the open eye template, the system determines that the eye has closed. The detected eye direction and blinks can then be used to control applications such as spelling programs or games. Prior work on this system focused on eye direction detection. With the addition of blink detection, three keyboard or mouse events can now be simulated. Experiments with EyeKeys have shown that it is an easily-used computer input and control device for able-bodied people and has the potential to become a practical tool for people with severe paralysis.