Seminar & Events Bulletin: Theoretical Computer Science

01-01-2013 to 06-30-2013

Friday, January 18, 2013	
10:30am-11:30am	**Theoretical Computer Science** -- Seth Pettie (U-M) *The Locality of Distributed Symmetry Breaking* -- 3941 BBB/CSE

Friday, January 25, 2013	
10:00am-11:00am	**Theoretical Computer Science** -- Travis Martin (U-M) *Characterizing Strategic Cascades on Networks* -- 411 West Hall

Friday, February 01, 2013	
10:00am-11:00am	**Theoretical Computer Science** -- Atri Rudra (Buffalo) *One algorithm to rule them all: One join query at a time* -- 411 West Hall

Friday, February 08, 2013	
10:00am-11:00am	**Theoretical Computer Science** -- Qi Cheng (University of Oklahoma) *On the Decodability of Primitive Reed-Solomon Codes* -- 411 West Hall

Friday, February 15, 2013	
10:30am-11:30am	**Theoretical Computer Science** -- Steve Lu () *Distributed Oblivious RAM for Secure Two-Party Computation* -- 3941 BBB/CSE

Friday, March 01, 2013	
10:00am-11:00am	**Theoretical Computer Science** -- Ely Porat (BIU/UM) *Sketching For Big Data Recommender Systems Using Fast Pseudo-Random Fingerprints* -- 411 West Hall

Friday, March 22, 2013	
10:00am-11:00am	**Theoretical Computer Science** -- Arnab Bhattacharyya (DIMACS) *Every locally characterized affine-invariant property is testable* -- 411 West Hall

Friday, April 05, 2013	
10:30am-11:30am	**Theoretical Computer Science** -- Aaron Snook (U-M) *An Optimal Lower Bound on the Number of Variables for Graph Identification* -- 3941 BBB/CSE

Friday, April 12, 2013	
10:00am-11:00am	**Theoretical Computer Science** -- Aram Harrow (MIT) *High-degree graphs cannot be used for a quantum PCP* -- 411 West Hall

Friday, April 19, 2013	
10:00am-11:00am	**Theoretical Computer Science** -- Mary Wootters (U-M) *What gaussian processes can do for you: applications of probability and geometry in theoretical computer science* -- 411 West Hall

Wednesday, May 01, 2013	
12:00pm-1:00pm	**Theoretical Computer Science** -- Valerie King (University of Victoria) *Dynamic Graph Connectivity in Polylogarithmic Worst Case Time* -- 4941 BBB/CSE

Seminar & Events Bulletin: Theoretical Computer Science

01-01-2013 to 06-30-2013

Friday, May 03, 2013	
10:30am-11:30am	**Theoretical Computer Science** -- Jared Saia (University of New Mexico) *Byzantine Agreement in Polynomial Expected Time* -- 3941 BBB/CSE

Seminar & Events Bulletin: Theoretical Computer Science

01-01-2013 to 06-30-2013

Abstracts

Theoretical Computer Science
Friday, January 18, 2013, 10:30am-11:30am
3941 BBB/CSE
Seth Pettie (U-M)
The Locality of Distributed Symmetry Breaking

We present new methods for solving several classical symmetry breaking tasks in distributed networks, such as finding maximal independent sets, maximal matchings, and vertex-colorings. This is joint work with Leonid Barenboim, Michael Elkin, and Johannes Schneider. An extended abstract appeared in FOCS 2012. PDF available at http://web.eecs.umich.edu/~pettie/papers/Symmetry-Breaking.pdf.

Theoretical Computer Science
Friday, January 25, 2013, 10:00am-11:00am
411 West Hall
Travis Martin (U-M)
Characterizing Strategic Cascades on Networks

I will present an in-progress project using game theory to capture a particular network process: cascades. Examples of processes which can be modeled as cascades are product adoption and opinion formation.

All current game theoretic cascade models use agents of limited strategic ability in order to simplify game dynamics. My work investigates the impact of this simplifying assumption by providing bounds on game behavior under with fully strategic agents. Due to the in-progress nature of this work, discussion and understanding will be heavily emphasized.

Current work and a more detailed abstract can be found at:
http://www-personal.umich.edu/~travisbm/publications/travis-prelim.pdf

Theoretical Computer Science
Friday, February 01, 2013, 10:00am-11:00am
411 West Hall
Atri Rudra (Buffalo)
One algorithm to rule them all: One join query at a time

We present a recent algorithm (PODS 2012) that is the first provably optimal (worst-case) algorithm to compute database joins.

As a special case, we show that this join algorithm implies (i) The first algorithmic versions of some well-known geometric inequalities due to Loomis and Whitney (and their generalizations by Bollobas and Thomason); (ii) Algorithmically list recoverable codes that work with parameters that no known algorithmic list recovery result work with (e.g. those based on the Reed-Solomon codes) and an application of this result in designing sublinear time decodable compressed sensing schemes; (iii) Worst-case optimal algorithm to list all occurrences of any fixed hypergraph H in a given large hypergraph G.

We believe that this algorithm should find many more applications. (If time permits, I'll also mention some followup work on instance optimal join algorithms.)

This talk will focus on (i) and (ii) and is based on joint works with Gilbert, Ngo, Nguyen, Porat, Re and Strauss.

Bio: Atri Rudra is an Assistant Professor of Computer Science and Engineering at University at Buffalo, State University of New York, Buffalo. Atri received his Bachelor's degree from Indian Institute of Technology, Kharagpur, India in 2000 and his Ph.D. from University of Washington in 2007. From 2000-2002, he was a Research Staff Member at IBM India Research Lab, New Delhi, India.

His research interests lie in theoretical computer science and in particular, theory of error-correcting codes, data stream and sub-linear algorithms, database algorithms, computational complexity, finite field theory and applications. He is a recipient of an NSF CAREER award (2009), HP Labs Innovation Research Award (2010), ESA best paper award (2010), UB Exceptional Scholars - Young Investigator award (2011) and PODS best paper award (2012).

Seminar & Events Bulletin: Theoretical Computer Science

01-01-2013 to 06-30-2013

Theoretical Computer Science
Friday, February 08, 2013, 10:00am-11:00am
411 West Hall
Qi Cheng (University of Oklahoma)
On the Decodability of Primitive Reed-Solomon Codes

Reed-Solomon codes are (list-)decodable up to the Johnson-Guruswami-Sudan bound. No polynomial time decoding algorithm is known when number of errors is larger than the JGS bound. The maximum likely-hood decoding of generalized Reed-Solomon codes is NP-hard, but it appears hard to establish complexity results for the primitive Reed-Solomon codes. In this talk, I will present several results on this problem. I will also talk about the deterministic construction of small Hamming balls containing many Reed-Solomon codewords.

Theoretical Computer Science
Friday, February 15, 2013, 10:30am-11:30am
3941 BBB/CSE
Steve Lu ()
Distributed Oblivious RAM for Secure Two-Party Computation

We present a new method for secure two-party Random Access Memory (RAM) program computation that does not require taking a program and first turning it into a circuit. The method achieves logarithmic overhead compared to an insecure program execution.

At the heart of our construction is a new Oblivious RAM protocol where a client interacts with two non-communicating servers. Our two-server Oblivious RAM for n reads/writes requires O(n) memory for the servers, O(1) memory for the client, and O(log n) amortized read/write overhead for data access. In our two-server model, we describe a new technique to bypass oblivious sorting which results in tiny constants and leads to a more practical Oblivious RAM protocol that compares favorably to the state-of-the-art single-server schemes.

Our two-server Oblivious RAM protocol leads to a novel application in the realm of secure two-party RAM program computation. We show that our Oblivious RAM construction can be composed with an extended version of the Ostrovsky-Shoup compiler to obtain a new method for secure two-party program computation with lower overhead than all existing constructions.

Joint work with Rafail Ostrovsky.

Seminar & Events Bulletin: Theoretical Computer Science

01-01-2013 to 06-30-2013

Theoretical Computer Science
Friday, March 01, 2013, 10:00am-11:00am
411 West Hall
Ely Porat (BIU/UM)

Sketching For Big Data Recommender Systems Using Fast Pseudo-Random Fingerprints

A key building block for collaborative filtering recommender systems is finding users with similar consumption patterns. Given access to the full data regarding the items consumed by each user, one can directly compute the similarity between any two users. However, for massive recommender systems such a naive approach requires a high running time and may be intractable in terms of the space required to store the full data. One way to overcome this is using sketching, a technique that represents massive datasets concisely, while still allowing calculating properties of these datasets. Sketching methods maintain very short fingerprints of the item sets of users, which allow approximately computing the similarity between sets of different users. The state of the art sketch has a very low space complexity, and a recent technique shows how to exponentially speed up the computation time involved in building the fingerprints. Unfortunately, these methods are incompatible, forcing a choice between low running time or a small sketch size. We propose an alternative sketching approach, which achieves both a low space complexity similar to that of [22] and a low time complexity similar to [14]. We empirically evaluate our algorithm using the Netflix dataset.We analyze the running time and the sketch size of our approach and compare them to alternatives. Further, we show that in practice the accuracy achieved by our approach is even better than the accuracy guaranteed by the theoretical bounds, so it suffices to use even shorter fingerprints to obtain high quality results.

Theoretical Computer Science
Friday, March 22, 2013, 10:00am-11:00am
411 West Hall
Arnab Bhattacharyya (DIMACS)
Every locally characterized affine-invariant property is testable

Let F = F_p for any fixed prime p >= 2. An affine-invariant property is a property of functions on F^n that is closed under taking affine transformations of the domain. We prove that all affine-invariant property having local characterizations are testable. In fact, we show a proximity-oblivious test for any such property P, meaning that there is a test that, given an input function f, makes a constant number of queries to f, always accepts if f satisfies P, and rejects with positive probability if the distance between f and P is nonzero. More generally, we show that any affine-invariant property that is closed under taking restrictions to subspaces and has bounded complexity is testable.

We also prove that any property that can be described as the property of decomposing into a known structure of low-degree polynomials is locally characterized and is, hence, testable. For example, whether a function is a product of two degree-d polynomials, whether a function splits into a product of d linear polynomials, and whether a function has low rank are all examples of degree-structural properties and are therefore locally characterized.

Our results depend on a new Gowers inverse theorem by Tao and Ziegler for low characteristic fields that decomposes any polynomial with large Gowers norm into a function of low-degree non-classical polynomials. We establish a new equidistribution result for high rank non-classical polynomials that drives the proofs of both the testability results and the local characterization of degree-structural properties.

Joint work with Eldar Fischer, Hamed Hatami, Pooya Hatami, and Shachar Lovett.

Seminar & Events Bulletin: Theoretical Computer Science

01-01-2013 to 06-30-2013

Theoretical Computer Science
Friday, April 05, 2013, 10:30am-11:30am
3941 BBB/CSE
Aaron Snook (U-M)
An Optimal Lower Bound on the Number of Variables for Graph Identification

In this paper we show that Omega(n) variables are needed for first-order logic with counting to identify graphs on n vertices. The k-variable language with counting is equivalent to the (k-1)-dimensional Weisfeiler-Lehman method. We thus settle a long-standing open problem. Previously it was an open question whether or not 4 variables suffice. Our lower bound remains true over a set of graphs of color class size 4. This contrasts sharply with the fact that 3 variables suffice to identify all graphs of color class size 3, and 2 variables suffice to identify almost all graphs. Our lower bound is optimal up to multiplication by a constant because n variables obviously suffice to identify graphs on n vertices.

Theoretical Computer Science
Friday, April 12, 2013, 10:00am-11:00am
411 West Hall
Aram Harrow (MIT)
High-degree graphs cannot be used for a quantum PCP

One variant of the quantum PCP conjecture states that it is QMA-complete to estimate the ground-state energy of a Hamiltonian with n qubits up to an error proportional to the total number of interacting pairs of qubits in the system. Since this generalizes classical 2-CSPs, this problem is at least NP-hard. We prove that the ground-state energy of 2-local Hamiltonians on D-regular graphs can be approximated in NP with additive error inverse-polynomial in D. Thus, if a quantum PCP theorem were to be true, it would need to make use of constant-degree graphs. The proof is based on information-theoretic techniques introduced by Raghavendra and Tan in arXiv:1110.1064.

Similar techniques also yields a PTAS for Hamiltonians on dense hypergraphs, planar graphs and highly expanding graphs. This last result in fact makes use of an application of the Lasserre SDP hierarchy to the quantum Hamiltonian problem, which generalizes its application to classical CSPs.

Based on joint work with Fernando Brandao.

Seminar & Events Bulletin: Theoretical Computer Science

01-01-2013 to 06-30-2013

Theoretical Computer Science
Friday, April 19, 2013, 10:00am-11:00am
411 West Hall
Mary Wootters (U-M)

What gaussian processes can do for you: applications of probability and geometry in theoretical computer science

In this expository talk, I'll give an introduction to gaussian processes and how they can be useful in theoretical computer science. In addition to discussing how to use this toolkit, I will highlight some recent applications in TCS, including results in compressed sensing, coding theory, and graph theory.

Theoretical Computer Science
Wednesday, May 01, 2013, 12:00pm-1:00pm
4941 BBB/CSE
Valerie King (University of Victoria)

Dynamic Graph Connectivity in Polylogarithmic Worst Case Time

The dynamic graph connectivity problem is the following: given a graph on a fixed set of n nodes, design a data structure to process an online sequence of updates in the form of edge insertions and deletions, and queries of the form q(a,b): "Is there a path between nodes a and b?" While data structures for this problem with polylogarithmic *amortized* time per operation have been known since the mid-1990's, these data structures have Theta(n) worst case time. In fact, no previously known solution has worst case time per operation which is o(sqrt n).

In this talk I'll explain a solution with worst case times O(log^4 n) per edge insertion, O(log^5 n) per edge deletion, and O(log n/loglog n) per query. The answer to each query is correct if the answer is "yes" and is correct with high probability if the answer is "no." The data structure is based on a simple novel idea which can be used to quickly identify an edge in a cutset.

This work is joint with Bruce Kapron and Ben Mountjoy.

Seminar & Events Bulletin: Theoretical Computer Science

01-01-2013 to 06-30-2013

Theoretical Computer Science
Friday, May 03, 2013, 10:30am-11:30am
3941 BBB/CSE
Jared Saia (University of New Mexico)
Byzantine Agreement in Polynomial Expected Time

How can we build a reliable system out of unreliable parts? Byzantine agreement is fundamental to addressing this question. The Byzantine agreement problem is to devise an algorithm so that n agents, each with a private input can agree on a single common output that is equal to some agent's input. In the classic Byzantine agreement problem, communication is via asynchronous message-passing and the adversary is adaptive with full information. In particular, the adversary can adaptively determine which processors to corrupt and what strategy these processors should use as the algorithm proceeds; the scheduling of the delivery of messages is set by the adversary, so that the delays are unpredictable to the algorithm; and the adversary knows the states of all processors at any time, and is assumed to be computationally unbounded. Such an adversary is also known as strong.

We present a polynomial expected time algorithm to solve asynchronous Byzantine Agreement with a strong adversary that controls up to a constant fraction of the processors. This is the first improvement in running time for this problem since Ben-Or's exponential expected time solution in 1983. Our algorithm is designed so that in order to thwart it, corrupted agents must engage in statistically deviant behavior that is detectable by individual agents. This suggests a new paradigm for secure distributed computing: the design of algorithms that force an attacker into behavior that is statistically deviant in a way that is computationally detectable.