Table of Contents

Survey talks (60 minutes each)

Shorter talks (40 minutes each)

Surveys

Joshua Brody: Property Testing Lower Bounds via Communication Complexity

Abstract: A few years ago, Blais, Brody, and Matulef introduced a method for proving property testing lower bounds via communication complexity [BBM12]. Since its introduction, this technique has become a standard tool for proving lower bounds for testing properties of functions. In this talk, I will start by introducing the property testing model. Then, I'll develop the connection to communication complexity and provide some examples of testing lower bounds achieved via communication complexity. The bulk of the talk will survey open problems in property testing. I assume the audience is familiar with the basics of communication complexity, but no prior knowledge of property testing is needed to follow the talk.

Arkadev Chattopadhyay: The state-of-affairs in the Number-on-Forehead model (a biased survey)

Abstract: Chandra, Furst and Lipton introduced the Number-on-Forehead (NOF) model of multiparty communication more than three decades ago. While it is perhaps not the most "intuitive" generalization of the classical two-party communication model, it and some of its variants continue to challenge researchers as it captures apparent communication bottlenecks in diverse computational situations like circuit complexity, pseudo-random generators, branching programs, proof complexity and data-structures.

In this talk, we will survey some of the challenges and some of the recent progress made in understanding the NOF model.

Hartmut Klauck: Quantum Communication Complexity

In this talk I will try to survey the field of quantum communication complexity with an emphasis on identifying some open problems that may not be widely known.

Kasper Larsen: Communication Complexity and Data Structures

Abstract: In this talk I will survey the use of communication complexity tools in proving data structure lower bounds. The talk has two halves, one focussing on static data structures and one on dynamic data structures.

In a static data structure problem, an input data set is given and must be preprocessed such that queries can be answered efficiently. The first application of communication complexity in proving lower bounds for static data structures is due to Miltersen et al. The basic idea is to consider a communication game in which Alice receives a query and Bob the input data set. If a data structure answers queries using only t memory lookups into a memory of S words, each containing w bits, then the communication game has a protocol in which Alice sends t log S bits and Bob sends t w bits, in which Alice and Bob simply simulate the query algorithm of the data structure. Proving lower bounds on the communication between Alice and Bob then yields lower bounds on the space S and query time t of the data structure. This technique was later refined by Patrascu and Thorup by letting Alice hold not only one query, but k queries. The new goal is to answer all of Alice’s queries on Bob one input data set. The crucial observation is that Alice can simulate all k queries in parallel and thereby ask for a subset of k memory words simultaneously. This costs roughly k log(S/k) bits compared to log S bits per memory lookup and thus a more efficient protocol is obtained. Again, data structure lower bounds are obtained by bounding the communication in this new setup. Finally I will briefly sketch how the cell probe model, normally underlying data structure lower bounds proofs, can be interpreted as a communication game in which Bob is just a table of S precomputed replies to S different messages that Alice can send. It seems that we have to take this extra structure into account if we are to advance the state-of-the-art in proving data structure lower bounds.

In a dynamic data structure problem, the goal is to maintain an underlying data set subject to updates and queries. As for static data structures, the highest obtained query time lower bounds are only polylogarithmic. For this part of the talk, I will focus on a three-party communication game defined by Patrascu. In this communication game, we have three players Alice, Bob and Carmen. Each of the players have an input on their forehead and thus can see only the inputs given to the other players. The input of Alice is an integer i in [k], Bob has k sets S_1,…,S_k \subset [n] and Carmen has one set T \subset [n]. The game rule is as follows: First Alice sends one message of length o(k) secretly to Bob. Following that, Bob and Carmen engage in a standard two-way communication game with the goal of answering whether set S_i intersects T. If a lower bound of n^c for any constant c>0 can be proved for the communication between Bob and Carmen, then polynomial data structure lower bounds follow. The underlying intuition for why they would need much communication is that Alice’s message to Bob holds much less than 1 bit of information about each S_i on average, and thus should not help. However, as shown by Chattopadhyay et al, this intuition is not always correct and that the exact behaviour in this advice model is much more intriguing.

Rotem Oshman: The Role of Communication Complexity in Distributed Computing

In distributed systems, the cost of communicating between the participating devices often dwarfs the cost of local computation on the individual devices. Accordingly, when we model a distributed system, we usually ignore local computation and charge only for communication. Traditional cost measures have included the total number of messages sent, or the number of rounds of communication, without restricting the number of bits in a message or a round. However, recently there has been a lot of interest in quantifying the number of bits that need to be sent to solve various tasks, and the interaction between this cost and the round complexity of the task. In this talk I will survey some of the recent work incorporating communication costs into distributed computing models, give some examples of lower bounds, and discuss open problems.

Omri Weinstein: Interactive Information Complexity

Abstract:
In the late 1940's Shannon introduced his information theory in order to understand the one way data transmission problem. Shannon's work revealed the intimate connection between information and communication, namely, that the amortized transmission cost of a random message is equal to the amount of information it contains. Since then information theory has been developed and applied in many different directions, and became the primary tool for analyzing communication problems.

Classical information theory, however, does not readily lend itself to interactive setups such as the Communication Complexity model, where two (or more) players must engage in a multi-round conversation in order to accomplish some desirable task. Our main quantity of interest is the Information Complexity of a problem, which informally captures the average amount of information the players need to disclose each other about their inputs in order to solve the problem. Motivated by many applications (in Communication Complexity, Privacy, Streaming and Circuit Complexity to mention a few), this field of research attempts to extend Shannon's theory, develop the right tools, and understand how information behaves in interactive setups. Indeed, the simple yet powerful benefits of this relaxed complexity measure (such as the chain rule and the data processing inequality) have had a profound impact on communication and circuit complexity lower bounds and hardness amplification over the past few years.

In this expository talk I will lay down the framework and main concepts which will be used in future related talks throughout the workshop. I will also describe some of the cornerstone results in the field, and how information complexity helped us make progress on some of the major open problems in TCS, providing additional evidence that information theory is the "right" tool for studying communication complexity.

David Woodruff: Lower Bounds for Data Streams

Abstract: I'll define a few data stream models and discuss techniques for proving lower bounds in them, giving a few examples obtained by looking at different models of communication complexity. The majority of the talk will be on surveying the open questions in this direction.

Talks

(authors in alphabetical order by last name)

Arkadev Chattopadhyay: A little advice can be very helpful

Abstract: Proving super-polylogarithmic lower bounds for dynamic data structures has remained an open problem. Patrascu proposed a new approach for breaking this barrier via a two player communication model in which one player gets private advice at the beginning of the protocol.

He gave reductions from the problem of solving an asymmetric version of set-disjointness in his model to a diverse collection of natural dynamic data structure problems in the cell probe model. He also conjectured that, for any hard problem in the standard two-party communication model, the asymmetric version of the problem is hard in his model, provided not too much advice is given.

In this paper, we prove several surprising results about his model. In particular, this disproves one of the conjectures of Patrascu. Additionally, we prove lower bounds for restricted protocols in Patrascu's model. These are now known to yield lower bounds for restricted data-structures.

Joint with J. Edmonds, F. Ellen and T. Pitassi.

Ankit Garg: Small value parallel repetition for general games.

Abstract: We prove a parallel repetition theorem for general games with value tending to 0. Previously Dinur and Steurer proved such a theorem for the special case of projection games. We use information theoretic techniques in our proof. Our proofs also extend to the high value regime (value close to 1) and provide alternate proofs for the parallel repetition theorems of Holenstein and Rao for general and projection games respectively. This is joint work with Mark Braverman.

Sudipto Guha: Linear Sketches and Graph Algorithms

Abstract: Linear Sketches are one of the most versatile and powerful tools in the development of sublinear algorithms. In this talk we consider how linear sketches can be used to develop graph algorithms in different models, in particular, semistreaming, adaptive sketching, and mapreduce models. In all these models the tradeoff between the number of rounds of computation and the space usage, is a natural measure of efficiency of different algorithms.

We show that linear sketches can be used to ``compress'' the number of rounds in a very natural way. Specifically, we reexamine the Ahn-Guha-McGregor SODA 2012 algorithm in this light and extend the intuition to develop new algorithms for maximum matching. Interestingly analysis ideas force us to relate the number of rounds of a sublinear algorithm to the number of rounds of Dantzig-Wolfe decompositions used to (approximately) solve linear programs. We expect this connection to linear programs can be used to develop tradeoffs between approximation, space usage and the number of rounds several graph problems.

Venkatesan Guruswami: Communication with Imperfectly Shared Randomness

The communication complexity of many fundamental problems reduces greatly when the communicating parties share randomness that is independent of the inputs to the communication task. Natural communication processes, however, often involve large amounts of shared correlations among the communicating players, but rarely allow for perfect sharing of randomness. Can the communication complexity benefit from shared correlations as well as it does from shared randomness? This question was recently studied mainly in the context of simultaneous communication by Bavarian et al. (ICALP 2014). In this work we study this question in the setting of one-way communication. We investigate fundamental problems such as compression, and agreement distillation (extracting perfectly shared randomness) in this context. We also show general connections between one-way communication complexity in the setting of perfectly shared randomness versus the setting of imperfectly shared randomness. In particular, we show that every problem with ``constant'' one-way communication complexity of k

bits in the former setting has a one-way protocol with 2^k bits in the latter setting. Our main technical result is a matching lower bound showing this exponential price in the simulation is best possible. The lower bound builds on the intuition that communication with imperfectly shared randomness needs to be less sensitive to inputs than communication with perfectly shared randomness. The formal proof invokes results about the small-set expansion of the noisy hypercube and an invariance principle to convert this intuition to a proof, thus giving a new application domain for these fundamental results.

Joint work with Clement Canonne, Raghu Meka, and Madhu Sudan.

Gillat Kol: Exponential Separation of Information and Communication

Abstract: We show an exponential gap between communication complexity and information complexity, by giving an explicit example for a communication task, with information complexity O(k), and distributional communication complexity Omega(2^k). This shows that a communication protocol cannot always be compressed to its internal information. By a result of Braverman, our gap is the largest possible. By a result of Braverman and Rao, our example shows a gap between communication complexity and amortized communication complexity, implying that a tight direct sum result for distributional communication complexity cannot hold.

Joint work with Anat Ganor and Ran Raz.

Frederic Magniez: Unidirectional Input/Output Streaming Complexity of Reversal and Sorting

joint work with Nathanaël François (U. Paris Diderot), Rahul Jain (National University of Singapore)

Abstract:
We consider unidirectional data streams with restricted access, such as read-only and write-only streams. We give tight bounds for the complexity of reversing a stream of length n in several of the possible models. In the read-only and write-only model, we

show that p-pass algorithms need memory space Θ(n/p). But if either the output stream or the input stream is read-write, then the complexity falls to Θ(n/p^2). It becomes polylog(n) if p=O(logn) and both streams are read-write.
We also quickly present the complexity of sorting a stream.

Jelani Nelson: Time lower bounds for nonadaptive turnstile streaming algorithms

Abstract: We say a turnstile streaming algorithm is "non-adaptive" if the memory cells it probes when receiving an update to coordinate x_i only depend on i as well as random coins flipped before seeing the stream (e.g. to determine hash functions). All known turnstile streaming algorithms are non-adaptive -- in fact, they are linear sketches.

We give the first ever space/update time tradeoff lower bounds that hold against non-adaptive turnstile streaming algorithms. Our lower bounds hold against classically studied problems such as heavy hitters, point query, moment estimation, and entropy estimation. In the case of deterministic L1 point query, known algorithms show that our lower bounds are tight for constant error parameter epsilon.

Manuscript online at http://arxiv.org/abs/1407.2151. Joint work with Kasper Green Larsen and Nguyễn Lê Huy.

Krzysztof Onak: Massive Parallel Communication and Geometric Graph Problems

Abstract: I will discuss the complexity of geographic graph problems in the Massive Parallel Communication model (MPC), which is an attempt at formalizing modern parallel systems such as MapReduce. In particular, I will sketch both positive and negative results for problems such as Minimum Spanning Tree and Earth-Mover Distance. It is a great open question whether communication complexity can be used to prove strong unconditional lower bounds in this model.

Joint work with Alexandr Andoni, Aleksandar Nikolov, and Grigory Yaroslavtsev.

Justin Thaler: Stream Computation and Arthur-Merlin Communication

Abstract: We continue the study of streaming interactive proofs (SIPs). In this setting, a client (verifier) needs to process a massive stream of data, arriving online, but is unable to store even a small fraction of the data. It outsources the processing to a commercial cloud computing service (prover), but is unwilling to blindly trust answers returned by this service. Thus, the service must both supply the final answer and convince the verifier of its correctness.

Our primary objects of study are "barely interactive'" SIPs. Specifically, we show that constant-round SIPs are surprisingly powerful, by giving polylogarithmic cost two- and three-round protocols for several "query" problems, including Index, Nearest Neighbor Search, and Range Counting. This was thought to be impossible based on previous work.

On the other hand, in order to study the *limitations* of constant-round SIPs, we introduce a new hierarchy of communication models that we call "online Interactive Proofs" (OIPs). We give upper and lower bounds that (1) characterize every finite level of the OIP hierarchy in terms of previously-studied communication complexity classes, and (2) separate the first four levels of the hierarchy. Our study of OIPs reveals marked contrasts and some parallels with the classic Turing Machine theory of interactive proofs.

Joint work with Amit Chakrabarti, Graham Cormode, Andrew McGregor, and Suresh Venkatasubramanian.

Emanuele Viola: The communication complexity of group multiplication

Abstract: How much communication is required to determine whether a bunch of elements from a group G multiply to 1 (for various groups and various ways to distribute the elements to the players)?

We briefly discuss the abelian case, including some recent progress in http://www.ccs.neu.edu/home/viola/papers/ccsum.pdf

Then we focus on the non-abelian case, also motivated by applications in obfuscation/leakage-resilient cryptography. We premiere recent progress using "fancy" group theory (PSL, Gowers' trick, etc.), and raise several questions in communication and group theory.

Joint work in progress with Eric Miles

Thomas Watson: Zero-Information Protocols and Unambiguity in Arthur-Merlin Communication

Abstract: We study whether information complexity can be used to attack the long-standing open problem of proving lower bounds against Arthur–Merlin (AM) communication protocols. Our starting point is to show that—in contrast to plain randomized communication complexity— every boolean function admits an AM communication protocol where on each yes-input, the distribution of Merlin’s proof leaks no information about the input and moreover, this proof is unique for each outcome of Arthur’s randomness. We posit that these two properties of zero information leakage and unambiguity on yes-inputs are interesting in their own right and worthy of investigation as new avenues toward AM.

• Zero-information protocols (ZAM). Our basic ZAM protocol uses exponential communication for some functions, and this raises the question of whether more efficient protocols exist. We prove that all functions in the classical space-bounded complexity classes NL and □L have polynomial-communication ZAM protocols. We also prove that ZAM complexity is lower bounded by conondeterministic communication complexity.

• Unambiguous protocols (UAM). Our most technically substantial result is a $\Omega(n)$ lower bound on the UAM complexity of the NP-complete set-intersection function; the proof uses information complexity arguments in a new, indirect way and overcomes the “zero-information barrier” described above. We also prove that in general, UAM complexity is lower bounded by the classic discrepancy bound, and we give evidence that it is not generally lower bounded by the classic corruption bound.

Joint work with Mika Goos, Toniann Pitassi, Thomas Watson

Omri Weinstein: An Interactive Information Odometer and Applications

Abstract: We introduce a novel technique which enables two players to maintain an estimate of the internal information cost of their conversation in an online fashion without revealing much extra information. We use this construction to obtain new results about communication complexity and information-theoretically secure computation.

As a first corollary, we prove a strong direct product theorem for communication complexity in terms of information complexity: If I bits of information are required for solving a single copy of f under \mu with probability 2/3, then any protocol attempting to solve n independent copies of f under \mu^n using o(n*I) communication, will succeed with probability only 2^{-Omega(n)}. This is the best one can hope for, as Braverman and Rao [BR11] previously showed that O(n*I) communication suffice to succeed with probability ~ (2/3)^n.

We then show how the information odometer can be used to achieve information-theoretic secure communication between two untrusted parties:
If the players' goal is to compute a function f(x,y), and f admits a protocol with information cost is I and communication cost C, then our odometer can be used to produce a "robust" protocol which: (i) Assuming both players are honest, computes f with high probability, and (ii) Even if one party is malicious, then for any k, the probability that the honest player reveals more than O(k(I+ log C)) bits of information to the other player is at most 2^{-Omega(k)}.

Finally, we outline a potential approach which uses our odometer as a proxy for braking state of the art interactive compression results: Any progress on interactive compression in the regime where I=O(\log C) will lead to new *general* compression results in all regimes.

Joint work with Mark Braverman

Grigory Yaroslavtsev: Round vs. Communication Tradeoffs

Abstract: Tradeoffs between the number of rounds and communication have recently emerged as a new direction of research in communication complexity. In this talk I will discuss recent results in this area. Similar questions have also attracted a lot of attention in the distributed computing community. Models for distributed computing such as MapReduce and Massive Parallel Communication (MPC) focus the goals of algorithm design on the number of rounds of communication. I will describe recent developments in this area with the emphasis on connection with communication complexity.

Amir Yehudayoff: Simplified lower bounds on the multiparty communication complexity of disjointness

Abstract: We shall discuss the number-on-forehead communication complexity of set disjointness for k parties on a universe of size n. We shall see that in the deterministic case it is at least order n/4^k, which nearly matches Grolmusz's upper bound. We shall also discuss a simplification of Sherstov's proof of an order \sqrt{n}/(k2^k) lower bound for the randomized communication complexity. Joint work with Anup Rao.