Workshop on

Recent Trends in Algorithms

July 26-28, 2023

School of Computer Sciences

NISER, Bhubaneswar

Speaker: Surender Baswana
Title: Vital edges and sensitivity oracles for (s,t)-mincuts
Abstract:

Let G=(V,E) be a directed graph with a designated source vertex s and a designated sink vertex t. Each edge of G has a positive capacity. We all know about the classical Ford-Fulkerson algorithm and how its elegant analysis establishes the strong duality between the value of the maximum flow from s to t and the minimum value of (s,t)-cut. An edge in G is said to be vital if its failure leads to decrease in the value of (s,t)-mincut. We shall interactively discuss some interesting combinatorial results on vital edges - their cardinality, cover, and structure. The sensitivity oracle for (s,t)-mincuts is a compact data structure that can efficiently report the (s,t)-mincut after the failure of any small set of edges from E. We shall interactively discuss various sensitivity oracles and their optimality.

The talk will also have interesting open problems related to vital edges and sensitivity oracles.

Prerequisite: It is strongly recommended that the audience has the knowledge of the design and analysis of the Ford-Fulkerson algorithm. This is the only prerequisite for this talk.

Speaker: Barna Saha
Title:
Abstract:

Speaker: Sayan Bhattacharya
Title: Recent Advances in Dynamic Matching
Abstract:

Consider a graph G = (V, E) that is undergoing a sequence of edge insertions/deletions. We want to design an algorithm that maintains a large matching in this dynamic graph G with small "update time". Here, the "update time" of an algorithm refers to the time it takes to handle the insertion/deletion of an edge in G. Ideally, we would like to ensure that the update time of our algorithm is polylogarithmic in the number of nodes in G. This problem has received considerable attention within the dynamic algorithms community in the past decade. In this talk, I will start with a survey of this topic, and then present an overview of some very recent developments which point to surprising connections between dynamic and sublinear algorithms.

Speaker: Sujoy Bhore
Title: Euclidean Steiner Spanner: Light and Sparse
Abstract:
Lightness and sparsity are two natural parameters for Euclidean $(1+\varepsilon)$-spanners. Classical results show that, when the dimension $d\in \mathbb{N}$ and $\varepsilon>0$ are constant, every set S of n points in d-space admits an $(1+\varepsilon)$-spanners with $O(n)$ edges and weight proportional to that of the Euclidean MST of S. In a recent breakthrough, Le and Solomon (FOCS'2019) established the precise dependencies on $\varepsilon>0$, for constant $d\in \mathbb{N}$, of the minimum lightness and sparsity of $(1+\varepsilon)$-spanners, and observed that Steiner points can substantially improve the lightness and sparsity of a $(1+\varepsilon)$-spanner. They gave lower bounds of $\tilde{O}(\varepsilon^{-(d+1)/2})$ for the minimum lightness in dimensions $d\geq 3$, and $\tilde{O}(\varepsilon^{-(d-1))/2})$ for the minimum sparsity in d-space for all $d\geq 1$. Subsequently, Le and Solomon (ESA'2020) constructed Steiner $(1+\varepsilon)$-spanners of lightness $O(\varepsilon^{-1}\log\Delta)$ in the plane, where $\Delta\in \Omega(\sqrt{n})$ is the \emph{spread} of S, defined as the ratio between the maximum and minimum distance between a pair of points. Moreover, they constructed spanners of lightness $\tilde{O}(\varepsilon^{-(d+1)/2})$ in dimensions $d\geq 3$.
We improved several bounds on the lightness and sparsity of Euclidean Steiner $(1+\varepsilon)$-spanners. We established lower bounds of $\Omega(\varepsilon^{-d/2})$ for the lightness and $\Omega(\varepsilon^{-(d-1)/2})$ for the sparsity of such spanners in Euclidean d-space for all constant $d\geq 2$. Our lower bound constructions generalize previous constructions by Le and Solomon, but the analysis substantially simplifies previous work, using new geometric insight, focusing on the directions of edges. Next, we showed that for every finite set of points in the plane and every $\varepsilon > 0$, there exists a Euclidean Steiner $(1+\varepsilon)$-spanner of lightness $O(\varepsilon^{-1})$; this matches the lower bound for $d=2$. We generalize the notion of shallow light trees, which may be of independent interest, and use directional spanners and a modified window partitioning scheme to achieve a tight weight analysis. In this talk, we will discuss trade-offs in the bounds of lightness and sparsity of Euclidean Steiner spanner.
*This is a joint work with Csaba Toth.

Speaker: Mrinal Kumar
Title: Fast multivariate multipoint evaluation
Abstract:
Multipoint evaluation is the computational task of evaluating a polynomial given as a list of coefficients at a given set of inputs. A straightforward algorithm for this problem is to just iteratively evaluate the polynomial at each of the inputs. The question of obtaining faster-than-naive (and ideally, close to linear time) algorithms for this problem is a natural and basic question in computational algebra. In addition to its own inherent interest, faster algorithms for multipoint evaluation are closely related to fast algorithms for other natural algebraic questions like polynomial factorization and modular composition.

In this talk, I will briefly survey the state of art for this problem, and discuss some recent improvements and applications.

Speaker: Aakanksha Agrawal
Title: Computing Square Colorings
Abstract:
A square of a graph G is the graph obtained from it by adding edges between every pair of vertices that are at distance at most 2 in it. The Square Coloring problem takes as input a graph G and a number q, and the objective is to check if the square of G can be properly colored using at most q colors. We remark that although G may have various nice properties like smaller treewidth, the same may not hold for its square. Thus, in a way, by studying the Square Coloring, we can exploit the structural simplicity of G in obtaining coloring of its square.

The main focus of this talk will be to understand the (parameterized) complexity of Square Coloring. We will see an algorithm for the problem with a somewhat unusual running time that we show is optimal under the Exponential Time Hypothesis. We will also look at a sketch of a subexponential algorithm for the problem when the input graph is planar.

Speaker: Raghu Meka
Title: Strong bounds for three-term progressions.
Abstract:
Suppose you have a set S of integers from {1,2,...,N} that contains at least N / C elements. Then for large enough N, must S contain three equally spaced numbers (i.e., a 3-term arithmetic progression)?
In 1953, Roth showed this is the case when C is roughly (log log N). Behrend in 1946 showed that C can be at most exp(sqrt(log N)). Since then, the problem has been a cornerstone of the area of additive combinatorics. Following a series of remarkable results, a celebrated paper from 2020 due to Bloom and Sisask improved the lower bound on C to C = (log N)^(1+c) for some constant c > 0.
This talk will describe a new work showing that C can be as big as exp((log N)^0.08), thus getting closer to Behrend's construction. Based on joint work with Zander Kelley.

Speaker: Prajakta Nimbhorkar
Title: Matchings with Fairness Constraints
Abstract:
Matchings, with as well as without preferences, form a core of many applications. However, various constraints might be imposed to ensure diversity and fairness among matched individuals. The complexity of finding a matching that meets all the constraints depends on the structure of the constraints. We will see some recent algorithmic results and future directions in this area.

Speaker: Sutanu Gayen
Title: FPRAS for computing the total variation distance in large dimensions.
Abstract:
In this talk, we will look at the problem of estimating distance between two high-dimensional distributions. Specifically, we will look at the problem of multiplicative approximation of the total variation (TV) distance. Distance estimation is equivalent to testing whether the two distributions are close or far in TV distance which appears in machine learning contexts in practice. While additive approximation of the TV distance via Monte Carlo sampling was well-known, we initiate a thorough study of the multiplicative approximation problem.

I will start by showing the exact computation of the TV distance between two product distributions is hard, as are several variants of this problem. Then, I will discuss an algorithm for multiplicative approximation for certain special cases by reducing it to the counting knapsack problem. Finally, we will see a randomized algorithm for multiplicatively approximating the TV distance between two Bayes nets of bounded tree-width. The later algorithm employs a novel technique based on partial couplings, which might be of independent interest.

Based on joint work with Arnab Bhattacharyya, Kuldeep S. Meel, Dimitrios Myrisiotis, A. Pavan, and N. V. Vinodchandran.

Speaker: Arijit Ghosh
Title: Testing of Index-Invariant Properties in the Huge Object Model
Abstract:
The study of distribution testing has become ubiquitous in the area of property testing, both for its theoretical appeal, as well as for its applications in other fields of Computer Science. The original distribution testing model relies on samples drawn independently from the distribution to be tested. However, when testing distributions over the n-dimensional Hamming cube
{0,1}^n for a large n, even reading a few samples is infeasible. To address this, Goldreich and Ron [ITCS 2022] have defined a model called the huge object model, in which the samples may only be queried in a few places.
In this work, we initiate a study of a general class of properties in the huge object model, those that are invariant under a permutation of the indices of the vectors in {0,1}^n, while still not being necessarily fully symmetric as per the definition used in traditional distribution testing.
We prove that every index-invariant property satisfying a bounded VC-dimension restriction admits a property tester with a number of queries independent of n. To complement this result, we argue that satisfying only index-invariance or only a VC-dimension bound is insufficient to guarantee a tester whose query complexity is independent of n. Moreover, we prove that the dependency of sample and query complexities of our tester on the VC-dimension is tight. As a second part of this work, we address the question of the number of queries required for non-adaptive testing. We show that it can be at most quadratic in the number of queries required for an adaptive tester of index-invariant properties. This is in contrast with the tight exponential gap for

general non-index-invariant properties. Finally, we provide an index-invariant property for which the quadratic gap between adaptive and non-adaptive query complexities for testing is almost tight.

Speaker: Rohit Gurjar
Title: Greedy algorithms, matroids, and parallel complexity
Abstract:
There are many combinatorial optimization problems such as minimum weight spanning tree, scheduling unit length tasks, finding a basis for a set of vectors, which admit a natural greedy algorithm. Matroids are defined as an abstraction of the common property that enables a greedy algorithm. Optimization with two matroid constraints has a lot more interesting applications, for example, bipartite matching, finding two disjoint spanning trees etc. We will consider some interesting questions on these from a parallel complexity perspective. In particular, we will see how the search version and the decision version are equivalent for these problems, even in a parallel setting.

Speaker: Nithin Verma
Title: Improved sublinear algorithms for testing permutation freeness
Abstract:
Given a permutation pi of length k, an array A is pi-free if there are no k array values that, when considered from left to right, have the same relative order as that of the permutation. For example, the array A has is (2,1)-free if there are no two indices i < j such that A[i] > A[j] and the array is (1,3,2)-free if there are no three indices i < j < v such that A[j] > A[v] > A[i]. In particular, the set of (2,1)-free arrays are simply the set of all sorted arrays.

The problem of testing pi-freeness is to distinguish, with constant probability, arrays that are pi-free from arrays that need to be modified in at least a constant fraction of their values to be pi-free. This problem, which is a generalization of the well-studied monotonicity testing problem, was first studied systematically by Newman, Rabinovich, Rajendraprasad and Sohler (Random Structures and Algorithms; 2019). They proved that for all permutations pi of length at most 3, one can test for pi-freeness in polylog n many queries, where n is the array length. For arbitrary permutations of length k > 3, they also described a simple testing algorithm that makes O(n^{1-1/k}) queries. Most of the followup work has focused on improving the query complexity of testing pi-freeness for monotone pi.

In this talk, I will present an algorithm with query complexity O(n^{o(1)}) that tests pi-freeness for arbitrary permutations of constant length, which significantly improves the state of the art. I will also give an overview of the analysis that involves several combinatorial ideas.

Joint work with Ilan Newman.

Speaker: Roohani Sharma
Title:
Abstract: