Research Statement:
Network design, (Wireless) Networking, and Game Theory

MohammadTaghi Hajiaghayi[*]

"Please visit my website at *http://www.mit.edu/~hajiagha* that contains an *updated* publication list (to which I refer in this statement) and links to actual papers to download."

I am a computer scientist working in both theory and system aspects of network design in a way that combines tactics from my background in both algorithms and experimental sciences. In my research, I seek to design useful algorithms and mechanism, deploy them in real-world settings, and possibly perform extensive experiments to characterize and explain their behavior and performance. I believe algorithm design, which is often used for optimization goals, is the heart of almost all forms of modern computing systems and play the most important role in the fields of computer science, operation research and several other areas of science and engineering. In the same time, simulation on experimental data obtained in the process of measuring deployed systems reveals important information about the performance in the underlying environment. This is the reason that while I am doing research in algorithms I consider both theory and system sides. My goal is to push the frontiers of our algorithmic understanding of models and computational problems by investigating how varying the problem and the underlying model influences the optimal choice of the algorithm and how we can apply these algorithms and possibly implement them for real-world applications.

I have done research and made major advances in four fields of (theoretical/practical) algorithm design as follows:

- Network Design, Networking and Routing;
- Wireless and Sensor Network Design;
- Auction Design & Game Theory and connection to Machine Learning; and
- Planar Network Design.

Though these fields seem somehow distinct, the tools that I built, used, and learned in each field have empowered me to solve problems in other fields as well, and this is one of the main reasons that I am pursuing these fields of study simultaneously. My contributions to each one of these fields are deep and extensive. Below I highlight some of my main contributions in each of these four fields of study. Also, I mention some of the major unanswered questions on which I am currently working.

1 Network Design, Networking and Routing

Network Design and networking with its many variants is one of the most active research areas in computer science involving researchers from System, Networks, Algorithm Design, Graph Theory, Discrete Optimization, Game Theory and Information Theory. Especially mathematical modeling of networks plays a vital role in the understanding of computer and communication networks and provides insights into questions such as *allocation of network resources, analysis and effects of competitive and/or cooperative agents, Internet protocols, wireless network protocols, network dynamics, queuing systems performance optimization, and network traffic and topology*. These models shed light onto fundamental performance limits and trade-offs in practical scenario. In addition, new problems in this area are constantly propounded by practitioners working in various aspects of network design such as construction, routing and staged deployment. Furthermore, many new design paradigms such as ATM, Ad hoc and Wireless networking add rich new flavors to existing problems. My goal here is to focus on this active area of applications of algorithms in networking to understand current trends, identify understudied areas, and formulate new directions for further investigation. Below you find some of my major work in this area.

We consider approximation algorithms for **buy-at-bulk network design** problems and resolve a fundamental problem in network design which was open for almost 10 years (the paper appeared in FOCS'06). Buy-at-bulk network design problems arise in settings where economies of scale and/or the availability of capacity in discrete units result in concave or more generally sub-additive[1] cost functions on the edges (different edges might have different cost functions). This problem is fundamental in design of telecommunication networks. The typical scenario is that capacity (or bandwidth) on a link can be purchased in some discrete units $u_1 < u_2 < \ldots < u_r$ with costs $c_1 < c_2 < \ldots < c_r$ such that the cost per

[*]Computer Science Department, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213, U.S.A., E-mail: hajiagha@mit.edu

[1]A function f is sub-additive if $f(x) + f(y) \geq f(x + y)$ for $x, y \geq 0$

bandwidth decreases $c_1/u_1 > c_2/u_2 > ... > c_r/u_r$. The capacity units are sometimes referred to as cables or pipes. A basic problem that needs to be solved in this setting is the following: given a set of bandwidth demands, install sufficient capacity on the links of an underlying network topology so as to be able to route the demands. The goal is to minimize the total cost. Alternatively, we can think of the following scenario. We have two independent cost functions on the links of the network: a buy cost $b(e)$ and a rent cost $r(e)$. We are given a set of demands between some pairs of nodes. A feasible solution for the multicommodity buy-at-bulk instance is to buy some links and rent some other ones such that all the demands can be routed and the total cost is minimized; the cost of every bought edge is $b(e)$ and for rented edges it is $r(e)$ per unit of flow (bandwidth) over that link. This problem generalizes several classical problems, such as minimum cost flow to the settings that the cost of each edge is a subadditive monotone function. It is known that the problem is hard to approximate within a factor of $\Omega(\log^{1/4-\epsilon} n)$ for any $\epsilon > 0$. We give the **first poly-logarithmic approximation algorithm** for this fundamental challenging problem. Interestingly our algorithm is very simple, intuitive, and easy to implement possibly for real-world applications.

We further generalize our approach for this problem in several subsequent papers to the case that we have costs on the vertices (routers) instead of edges, to stochastic two-stage Steiner tree when inflation factors of diffract edges is different in two stages, and to the famous dial-a-ride problem that we want to find a minimum length path of a shuttle which carries several persons from their sources to their destination without violating the capacity k of the shuttle. On related topics, I have also published on classic k-forest, prize-collecting Steiner forest and k-Steiner buy-at-bulk problems.

Another fundamental concept in my line of research in this area is the role of **obliviousness** in network optimization and especially routing which has attracted lots of attention recently. The next-generation Internet will be orders of magnitude bigger in scale, connecting together nodes that are very different in their capabilities and requirements. Traditional algorithm design has tended to focus on problems with complete information over a homogeneous platform. However, the dramatic increase in scale and heterogeneity means that nodes can neither hope to obtain nor store and process full information. Therefore, it is important today to focus on algorithms and protocols that can operate obliviously, i.e., with limited knowledge. This research develops formal frameworks for quantifying the associated tradeoffs in oblivious network optimization and delivers infrastructure-class algorithms that can secure the foundations of the Internet of tomorrow. This area of research combines algorithm design and networking. Obtaining such oblivious solutions especially for network routing is a very challenging area on which currently I am working.

One of my very important work in this area is development of a framework to model **oblivious network design** and **(semi-)oblivious routing**. Consider the buy-at-bulk network design defined above in which all edges have uniform sub-additive cost function. Here we route each source-sink demand pair obliviously: when a terminal demand pair makes its routing decisions, it does not know the current flow on the edges of the network, nor the identity of the other pairs in the system. Moreover, it does not even know the identity of the function, merely knowing that is a sub-additive function of the total flow on the edge. Surprisingly we provide such an oblivious routing for each pair with $O(\log^2 n)$ competitive ratio with respect to an optimum solution which knows all demand pairs and their demands in advance and then reserve bandwidth on each edge. Here all our algorithms are simultaneously function-oblivious as well as traffic-oblivious and approximately minimizes the well-known measures of congestion and total cost of a network.

The above oblivious network design results work for undirected graphs. We further investigate oblivious routing on general directed graphs for measures of congestion and throughput (for the measure of total cost the problem is harder than notorious directed Steiner forest which is known to be hard to approximate within a polylogarithmic factor) which so far has led to many subsequent improvements and applications. In particular, in joint papers with Awerbuch, Kleinberg, Leighton and Räcke (SODA'05), we show that for directed networks even when there is only one sink in the network, the competitive ratio for both congestion and throughput is $\Omega(\sqrt{n})$ and indeed we can obtain essentially such competitive ratios for these cases. On the other hand, in joint work with Kim, Leighton and Räcke (STOC'05), we show that interestingly if we consider *oblivious routing with known demand distributions*, in which demands between all commodity pairs in the network are drawn i.i.d. from some distributions agreed in advance between the adversary and the algorithm, such a polylogarithmic competitive ratio is possible. We show that the lower bound on the competitive ratio is still polynomial if distributions are unknown to the algorithm or even when distributions are known to the algorithms but they are not independent. Most of the proofs of the lower bounds use algebraic/combinatorial graph structures and the proofs of the upper bounds use novel ideas for averaging the flow paths and/or the max-flow min-cut relation. Also, we consider node-capacitated undirected graphs for which we interestingly obtain similar upper/lower bounds to those in directed graphs.

We further investigate oblivious solutions for other routing problems such as traveling salesman problem (TSP); surprisingly we obtain an oblivious solution for TSP with competitive ratio $O(\log^2 n)$ which is a universal permutation that for any subset of demands visiting this demand set following our universal permutation gives a solution which is within a factor $O(\log^2 n)$ of the optimum tour for that demand set. In addition, we obtain the first non-constant lower bound $\Omega\left(\sqrt[6]{\frac{\log n}{\log\log n}}\right)$ for the n-by-n grid graph which refutes Bartholdi and Platzman conjecture (in JACM) from the late 1980's.

My other works in network routing (and not necessarily oblivious routing) include interval routing schemes (IRSs) that is a well-known space-efficient routing strategy for routing messages in distributed networks (Networks'02, TCS'04) and bandwidth sharing network design for multi-class traffic via simulation studies (on real-world networks in Bell-Labs), motivated by the increasing commercial interest in supporting voice and multimedia services over the IP network (INFOCOM'06).

2 Wireless and Sensor Networks

In recent years we had a great amount of research in wireless networks, especially wireless ad hoc and sensor networks. These works involve networking, routing, distributed algorithms, fault-tolerant and low-interference algorithms, mechanism design, computational geometry, and even combinatorics. However, there are still many challenges in wireless ad hoc networks. Due to the limited capability of processing power, storage, and energy supply, many conventional algorithms are too complicated to be implemented in wireless ad hoc networks. Some other algorithms did not take advantage of the geometry nature of the wireless networks. Additionally, most of the currently developed algorithms for wireless networks assumed a precise position of each wireless node, which is impossible practically. Majority of the algorithms developed in this area also assume all nodes have uniform transmission range. These algorithms will likely fail when nodes have disparity transmission ranges. Furthermore, the wireless ad hoc networks require efficient distributed algorithms with low computation complexity and low communication complexity. These algorithms are expected to take advantage of the geometry nature of the wireless ad hoc networks. Several fundamental questions should be answered: can we improve the performance of traditional distributed algorithms under wireless ad hoc networks? Does the position information of wireless nodes make difference in algorithms' performance? Much of the existing work in wireless ad hoc networking also assumes that each individual wireless node (possibly owned by selfish users) will follow prescribed protocols without deviation. However, each user may modify the behavior of an algorithm for self-interested reasons. How to achieve desired global system performances when individual nodes are selfish?

Below I address progress toward some of the above goals obtained in my current research.

My first work in this area (from two papers in WINET'06 and MOBICOM'03) pioneered the concept of power optimization in fault-tolerant topology control algorithms for wireless multi-hop sensor networks. More formally, here our goal is to find a k-connected spanning subgraph of a given graph which minimizes the power. The power of a spanning subgraph is the sum of the powers of the nodes, and the power of each node is the weight of the maximum edge attached to that node in the subgraph. Our work in this area, which presents efficient local heuristic algorithms and global practical distributed $O(k)$ approximation algorithms, got references from more than 60 papers in the wireless network community just in the past three years. In these papers, we also compare different approximation algorithms and heuristics via simulation. For large k, in a joint work with Kortsarz, Mirrokni and Nutov (IPCO'05), we also obtain a relatively simple (to implement) $O(\log^4 n)$ approximation algorithm. This line of research has been pursued by a recent joint work with Bredin, Demaine and Rus (MOBIHOC'05), in which we considered the problem of deploying or repairing a sensor network to provide simultaneously fault tolerance against node failures and high capacity through multi-path routing. We design, analyze and simulate the first constant-factor approximation algorithm that places an almost-minimum number of additional sensors to augment an existing network into a k-connected network, for any fixed k.

In closely related work to the above application and embeddings of networks (joint work with Badoiu, Demaine and Indyk in SoCG'04), we consider the problem of reconstructing the global geometry (i.e., assigning coordinates to the sensors in the plane or in space) of an ad hoc wireless sensor network from partial information about its geometry (e.g., distances, angles, and/or orientations). This global geometry can be used as a base for several other applications. We present several practical approximation algorithms for this problem, which originally arose in the SLAM (Scalable Location Aware Monitoring) project in MIT for sensor networks.

Motivated by applications mainly in wireless sensor networks and again closely related to above problems, we consider a general class of problems (in a recent SODA'07 paper) regarding minimizing movement. In general, these problems involve planning the coordinated motion of a large collection of objects with limited wireless capability (representing anything such as an autonomous robot swarm or firefighter) to achieve a global property of the network while minimizing the maximum or average movement. In particular, we consider the goals of achieving connectivity (undirected and directed), achieving connectivity between a given pair of vertices, achieving independence (a dispersion problem), and achieving a perfect matching (with applications to multicasting). These problems generally consist of a graph and a set of pebbles initially located on the vertices and we want to move the pebbles along the edges in order to reach a certain property, for example, connectivity, i.e., the induced subgraph formed by occupied pebbles is connected. This general family of movement problems encompasses an intriguing range of graph and geometric algorithms, with several real-world system applications and a surprising range of approximability. We could give distributed approximation algorithms and inapproximability results for this wide class of movement problems.

Wireless LAN administrators often have to deal with the problem of sporadic client congestion in popular locations within the network. Existing approaches that relieve congestion by balancing the traffic load are encumbered by the modifications that are required to both access points and clients. We propose *Cell Breathing*, a well known concept in cellular telephony, as a load balancing mechanism to handle client congestion in a wireless LAN. We develop power management algorithms for controlling the coverage of access points to handle dynamic changes in client workloads. We further incorporate hand-off costs and manufacturer specified power level constraints into our algorithms. Our approach does not require modification to clients or to the standard. It only changes the transmission power of beacon packets, and does not change the transmission power of data packets to avoid the effects of auto-rating. We analyze the worst-case bounds of the algorithms, and show they are either optimal or close to optimal. In addition, we evaluate our algorithms empirically using synthetic and real wireless LAN traces. Our results show that cell breathing significantly out-performs the commonly used fixed power scheme, and performs at par with sophisticated load balancing schemes that require changes to both the client and access points (The paper is a joint work with researchers in Microsoft Research and is to appear in IEEE Trans. on Mobile Computing).

Finally, we also consider the problem of power assignment and its relation to game theory and auction design (SODA'06), in which we want to assign powers to access points such that if each client chooses its "preferred" access point to which it connects, no access point exceeds its capacity. We show that there is a very close relation between this problem and *envy-free assignment* in auction design and we can use the ideas from mathematical economics to obtain power assignment for our practical scenario in Bell Labs and Microsoft Research .

3 Auction Design & Game Theory: connection to Machine Learning

Many of the key algorithmic challenges in the context of the internet require considering the objectives and interests of the different participants involved. These include problems ranging from pricing goods and resources, to improving search, to routing, and more generally to understanding how incentives of participants can be harnessed to improve the behavior of the overall system. As a result, Mechanism Design and Algorithmic Game Theory, which can be viewed as "incentive-aware algorithm design," have become an increasingly important part of algorithmic research in recent years. Along different lines, the area of Machine Learning has made continued progress in developing methods that can generalize from data, adapt to changing environments, and improve performance with experience, as well as progress in understanding of fundamental underlying issues. It is clear these capabilities will be critical as well in the context of such large decentralized systems. While these areas seem quite distinct, recent results in the area shows a strong relation between auction design and machine learning, and techniques from each seem well-poised to help with key problems of the other. My research interest here is to further develop these connections in order to produce powerful mechanisms for adaptive and networked environments, and improve the experience of users of the Web and internet. Currently we (joint with Avrim Blum, Manuel Blum, Michael Kearns, Tuomas Sandholm) **submitted a recent grant** on this topic. Below I highlight some of my previous work in his area.

Previous works in the area of auction design consider a static environment in which all bidders are present up front. If bidders instead arrive one at a time, then so long as they cannot game their timing information, we can view this as an online learning problem, and apply online learning techniques for auction problems that have appropriate structure . However, if bidders can misrepresent their timing information as well as their valuations, then this problem becomes much more difficult. Specifically, consider the problem of auction design for dynamic environments in which agents arrive and depart dynamically and in which goods are inherently temporal. Examples of such settings include the problem of WiFi allocation, sale of last-minute tickets, scheduling of scientific grid computing and projects such as MoteLab (at Harvard University) and PlanetLab (at Princeton University). In all these settings, a key challenge is to design **online auctions** that are incentive-compatible with respect to timing information, so users are not motivated to misrepresent their urgency or lack thereof. Here, even social-welfare optimization can be quite difficult. In particular, one of the most important techniques for designing truthful mechanisms (the Vickrey-Clarke-Groves (VCG) scheme) is inapplicable in most online problems because it requires computing an optimal allocation, which is generally impossible in the online setting. In our work (ACM EC'04 and ACM EC'05), motivated by auctions of digital goods, say in eBay or Amazon.com, we present several truthful mechanism which are constant competitive for revenue and/or efficiency. In addition we provide a characterization for the design of truthful online auctions, such that it is a dominant strategy equilibrium for bidders to reveal their true value for resources immediately upon arrival into a system. The auctions are online, in the sense that they make allocation decisions without knowledge of the future. In a setting without priors, we provide an e-competitive (for efficiency) truthful auction for a limited-supply unit-demand problem, drawing an analogy with the **classic secretary**

problem, a well-known concept in statistics and machine learning. We also show (in a model that slightly generalizes the assumption of independent valuations) that no mechanism can be better than 3/2- competitive (2-competitive) for revenue (efficiency). Our general approach considers a learning phase followed by an accepting phase, and is careful to handle incentive issues for agents that span the two phases.

We (in ACM EC'05 paper) also study the case of re-usable goods, such as processor time or wireless network access, which can be assigned to different agents at different times. Each agent is assumed to arrive and depart dynamically, and in the basic model requires the resource for one unit of time. In this work, we provide characterizations for the class of truthful online allocation rules and also present an online auction for unit-length jobs that achieves total value that is 2-competitive with the maximum offline value. We further prove that no truthful deterministic online mechanism can achieve a better competitive ratio. Finally, we generalize our model to settings with multiple re-usable goods and to agents with different job lengths.

Pursuing our work on relation between auction design and machine learning, we (in a joint work with Kleinberg and Sandholm) consider online auctions when we do not know exact number n of bidders in advance, which is well-motivated by applications in Ad-word auctions or Airline Tickets auctions in which we do not know the number of bidders in advance. Interestingly, we again find close relation between this concept and the classic **Prophet inequality**, a well-known concept in statistics. By exploiting and generalizing prophet inequality, we are able to design several truthful mechanisms for the case of unknown n. For the case in which we know the distribution of n, we exploit *automated mechanism design* to obtain truthful mechanisms.

In the same avenue of research on auction design and machine learning, in a joint work with Butler, Kleinberg and Leighton, we also consider **Hat guessing games** in which, in the typical case, each of n players tries to guess the color of the hat they are wearing by looking at the colors of the hats worn by some of the other players. One of the main applications of generalized hat guessing games is in the design of deterministic auction mechanisms. Another application is in gathering unknown information by exploiting some known information in a market.

Interestingly, some of my work above also shows fusion of ideas between *game theory and wireless network design* (e.g. pricing access to a WiFi port at Starbucks). This fusion has been further investigated in two joint papers with Bahl, Mirrokni, Qiu and Saberi and with Demaine, Feige and Salavatipour by considering power assignment in wireless networks, in which we want to assign powers to access points such that if each client connect to its "preferred" access point, no access point exceeds its capacity. We show that there is a close relation between this power assignment problem and the notion of **envy-free assignment** in game theory and we can use the ideas for this well-known concept to obtain power assignment for the access points. We also obtain the first almost-logarithmic hardness for envy-free pricing through a new similar hardness result for a maximization problem called the *unique coverage* problem. Concerning envy-free pricing, in a recent joint work with Blum, Balcan and Chan, we are considering envy-free pricing auctions with *negative prices*, in which in some cases we can sell our goods below their production costs and interestingly gain even more revenue.

Finally in a joint work with Jain (SODA'06), we also consider a new primal-dual scheme for prize-collecting generalized Steiner tree and k-forest. The novelty of this primal-dual scheme is trying to avoid (or at least simplify) cost-sharing issues when they arise during many such primal-dual schemes.

4 Planar Networks and Their Generalizations

Planar networks are graphs that can be drawn on the plane with no crossings. Such graphs naturally arise in application areas such as *image processing, logistics in road maps, high-speed fiber networks in the ground and VLSI*. In addressing fundamental optimization problems on graphs, if the input graphs can assumed to be planar, algorithms can be used that are faster and/or more accurate than algorithms that do not require this assumption. In particular, we (in joint work mainly with Demaine) introduced the newly developing theory of bidimensional graph problems in a series of 18 papers published e.g. in SODA/FOCS/STOC/JACM/SICOMP/SIDMA/TALG/Algorithmica. (The reader is referred to the survey by Demaine and Hajiaghayi[2] to see a full description of this theory.) This theory provides general techniques for designing efficient fixed-parameter algorithms[3] and approximation algorithms for NP-hard graph problems in broad classes of networks including planar graphs and more generally graphs with a few number of crossing. This research involves the design and analysis of such planarity-exploiting algorithms for fundamental optimization problems. Advances in this area could lead to improvements in methods for solving problems in road maps such as routing, network design,

[2]The survey was an invited presentation in the 12th International Symposium on Graph Drawing (GD), 2004, New York City, 2004. It can be accessed online at authors' homepages or directly at *http://www-math.mit.edu/~hajiagha/GD04.pdf*.

[3]A *fixed-parameter algorithm* is an algorithm for computing a parameter $P(G)$ of a graph G whose running time is $h(P(G))\, n^{O(1)}$ for some function h. (A *parameter* $P = P(G)$ is a function mapping graphs to nonnegative integers.) A typical function h for many fixed-parameter algorithms is $h(k) = 2^{O(k)}$.

and facility location, and for solving problems in image processing such as segmentation and tracking. In particular such graphs become particularly important recently since **according to several system talks in Workshop on Flexible Network Design 2005, almost all high-speed fiber networks in the ground are planar or they have a few number of crossings**. This theory applies to graph problems that are *bidimensional* in the sense that (1) the solution value for the $k \times k$ grid graph (and similar graphs) grows with k, typically as $\Omega(k^2)$, and (2) the solution value goes down when contracting edges and optionally when deleting edges. Examples of such problems include feedback vertex set, vertex cover, minimum maximal matching, face cover, clique transversal set, a series of vertex-removal parameters, dominating set, edge dominating set, R-dominating set, connected dominating set, connected edge dominating set, connected R-dominating set, and unweighted TSP tour (a walk in the graph visiting all vertices). *Lots of our algorithms here are easy to implement in practice and in addition can be used as simple heuristics even for general networks.*

Bidimensional problems have many structural properties; for example, any graph in an appropriate minor-closed class has *treewidth*, **a well-known concept in Artificial Intelligence** (roughly speaking, the measure of similarity of a graph to a tree) bounded above in terms of the problem's solution value, typically by the square root of that value. These properties lead to efficient—often subexponential—fixed-parameter algorithms, as well as polynomial-time approximation schemes, for many minor-closed graph classes. One type of minor-closed graph class of particular relevance has *bounded local treewidth*, in the sense that the treewidth of a graph is bounded above in terms of the diameter; indeed, as a **major result**, we show that such a bound is always at most linear (the previous bound was doubly exponential).

The bidimensionality theory unifies and improves several previous results. The theory is based on algorithmic and combinatorial extensions to parts of the seminal Robertson-Seymour Graph Minor Theory, in particular initiating a parallel theory of graph contractions. The foundation of this work is the topological theory of drawings of graphs on surfaces and our **major recent result** which says any graph excluding a fixed graph H as a minor, of treewidth w has an $\Omega(w) \times \Omega(w)$ grid graph as a minor (this is the tightest result up to constant factors). In this theory, we improve and generalize several results of others including Seymour, Robertson, Thomas, Alon, Eppstein, Bodlaender, Niedermeier, Reed, Lipton, Tarjan, Baker, Diestel, Thomassen, etc.—namely linearity of local treewidth (the previous bound was doubly exponential), linearity of the size of the grid minor in terms of treewidth, reproving the separator theorem, tight parameter-treewidth bounds, 2-approximation algorithms for coloring (in parallel to Hadwiger's Conjecture), partitioning (vertex set/edge set of) a graph into graphs of bounded treewidth, and PTASs for general *minimization* problems using our new separator theorem (which improves upon the classic Lipton-Tarjan and Alon-Seymour-Thomas theorems) in planar graphs and more generally minor closed graph families. (See the survey by Demaine and Hajiaghayi.) Some of the above results solved classic problems that were unsolved despite frequent study for over 15 years.

In this line of research, we (in a joint work with Feige and Lee in a selected paper in STOC'05) also develop the algorithmic theory of vertex separators, and its relation to the embeddings of certain metric spaces. We obtain an $O(\sqrt{\log \mathrm{opt}})$ approximation for min-ratio vertex cuts in general graphs, based on a new semidefinite relaxation of the problem, and a tight analysis of the integrality gap which is shown to be $\Theta(\sqrt{\log \mathrm{opt}})$. We also prove various approximate max-flow/min-vertex-cut theorems, which in particular give a constant-factor approximation for min-ratio vertex cuts in any excluded-minor family of graphs. Previously, this was known only for planar graphs, and for general excluded-minor families, the best-known ratio was $O(\log \mathrm{opt})$. Interestingly, we obtain tight bounds not only for general graphs, but also for other classes, e.g. planar graphs, for which corresponding results are not known in the edge case.

Our above results on vertex separator have a number of applications. We exhibit an $O(\sqrt{\log n})$ pseudo-approximation for finding balanced vertex separators in general graphs. Furthermore, **we obtain improved approximation ratios for treewidth (a very important concept both in theory and system):** *In any graph of treewidth k, we show how to find a tree decomposition of width at most $O(k\sqrt{\log k})$ (i.e., $O(\sqrt{\log opt})$-approximation)*, improving upon the previous $O(\log \mathrm{opt})$ approximation algorithms. For *planar graphs and more generally graphs excluding a fixed graph as a minor, we give an $O(1)$ approximation for the treewidth.* This result improves running times of many (practical) algorithms especially since treewidth often appears in the exponent of the running times. In addition via the bidimensionality theory, we can obtain the first polynomial-time approximation schemes for problems like minimum feedback vertex set and minimum connected dominating set in graphs excluding a fixed minor.

5 Conclusion

In summary, I study a variety of research areas (with both theory and system applications) in which algorithm design arise. Most of the areas still offer major unsolved problems (lots of them have important system applications) that would have considerable impact on the state-of-the-art of algorithm/network/mechanism design. Superficially, these areas seem to have little connection to each other. Our viewpoint reveals a little-known and surprisingly deep connection between all of these areas, which we believe is critical to solving these problems. In this sense, we are ideally poised to solve remaining open problems in these areas.