RESEARCH STATEMENT

Imran A. Pirwani

pirwani@cs.uiowa.edu

I am a researcher in Computer Science interested in designing and analyzing algorithms for problems in distributed computing and wireless networks. I am also interested in related fields such as graph theory, combinatorial optimization, computational complexity, and computational geometry.

Wireless networking lies at the cross-roads of theory and practice. While many problems encountered pose engineering challenges, these problems also are amenable to rigorous theoretical analysis and clean algorithmic solutions. Algorithmic study of these problems on "clean" models often provide significant insight into obtaining good real-life solutions. Much of my work has been on various geometric and combinatorial optimization problems in wireless networks. A wireless network consists of nodes in a terrain that are able to communicate with nearby nodes via radio broadcast and perform local computations based on information gathered from nearby nodes or the local environment. Nodes typically come equipped with limited battery supply. Common wireless network models have included *unit disk graphs* (UDG)[1], graphs in *doubling metric space*[2], *growth bounded graphs*[3], and *finite Euclidean metrics*. Some of my work has been driven by the need to extend the lifetime of wireless networks by: (i) carefully scheduling node activity [20], and (ii) communication using small band radio [12] to avoid interference. Motivated by the need to have location information for fast memory-less routing protocols, some of my recent work has addressed the problem of constructing virtual coordinates of a UDG from mere connectivity information [21]. Other work on geographic routing goes beyond the standard face routing protocols and is motivated by realistic graph models beyond UDGs that cannot be made planar and kept connected, simultaneously [17]. Most recently I have worked on clustering problems also motivated by challenges in wireless networking such as *minimum cost k–clustering* [10, 11]. In the following, I describe some of my research and a plan for the future.

Energy Conservation in Wireless Sensor Networks

Nodes in a wireless sensor network (WSN) typically come equipped with battery and from the point of deployment, this battery reserve becomes a valuable resource since it cannot be replenished ordinarily. Hence, maximizing the lifetime of the network by minimizing energy consumption is a fundamental challenge in WSNs. A standard approach for reducing energy consumption is to carefully schedule node activity. As has been observed in [6], whenever there are sufficiently many nodes in a region, only a small fraction of nodes need be active for forwarding messages, etc. The rest of the nodes can enter a *sleep* mode, thereby conserving energy without sacrificing coverage responsibility. Cycling through various disjoint subsets of such nodes keeps the network alive for longer than if all nodes were awake simultaneously. The problem of maximizing the number of nodes that are asleep at any given time while maintaining sufficient activity in the network is usually modeled as the problem of finding a small *dominating set* in the network. Finding a large partition of disjoint dominating sets becomes the problem of finding a large *domatic partition* in the underlying graph [18].

In our MobiHoc 2006 paper [20], we consider the problem of finding a large domatic partition on a UDG. We devise constant-factor approximation algorithms for a problem that is closely related to the domatic partition problem. We present three algorithms for successively weaker network models for wireless networks: UDGs with geometry, unit ball graphs in doubling metric spaces, and growth bounded graphs. We show, for example, that if access to geometry is available then a *2-domatic partition*[4] can be computed quickly whose size is at most a constant-factor worse than the size of an optimal domatic partition. While we do not obtain a constant-factor approximation to the domatic partition problem on UDGs, the solution is still quite relevant in WSNs. The implication of our result is that if the radio signal of each node of a dominating set is boosted by a small additional amount, then our solution is able to cover the sensors and its size is quite large relative to the size of the largest partition possible, if the signal hadn't been boosted.

---

[1] A UDG consists of a set of vertices in the Euclidean plane. Two vertices are connected iff their distance is at most 1.

[2] A metric space in which any ball of radius $2r$ can be cover with $O(1)$ balls of radius r.

[3] A graph is growth bounded if any ball of radius r contains at most $r^{O(1)}$ independent vertices

[4] Every dominating set is at most a 2-dominating set, where a 2-dominating set D is a subset of the vertices such that every vertex is at most distance 2 from some vertex in D.

Constructing Geometry from Topology in Wireless Networks

Multi-hop routing is a common service in WSNs, however, careless forwarding of messages needlessly drains valuable battery reserve. Knowledge of location information can dramatically improve the performance of routing algorithms because it allows the use of *geometric routing* techniques [4, 13, 15]. One technological solution for obtaining location information is to equip each node with a GPS receiver.

However, this solution seems too costly, currently. Recent work has suggested that for geometric routing schemes, having "real" coordinates is not necessary; having *virtual coordinates* suffices to ensure prompt and guaranteed routing [19, 23]. Using UDGs to model a WSN, Moscibroda et al. [14] consider the problem of computing *virtual coordinates* (an embedding in Euclidean plane) from only topological information. Since the coordinates are virtual, Moscibroda et al. define the measure of the *quality* of the embedding to be the ratio of the lengths of the longest edge to the shortest non-edge. An algorithm due to Vempala [24] gives an $O(\log^3 n)$ quality embedding for this problem. This solution depends upon using the *ellipsoid method* for solving a linear program of exponential size (in the number of constraints) to devise a solution.

In our work that appeared in ESA 2007 [21], we give a $O(\log^{2.5} n)$ quality for this *UDG realization* problem. While the quality is marginally better than Vempala's solution, the key feature of our solution is that it is completely combinatorial. The solution depends upon exploiting certain key properties of UDGs which enabled us to devise a first known constant-size clique partition of a neighborhood in a UDG without the aid of any geometry. This key structure is computed efficiently in a purely distributed manner.

Topology Control and Geographic Routing in Wireless Networks

I have also worked on ways to improve geographic routing in wireless networks that go beyond standard face routing protocols. This work is motivated by realistic graph models beyond UDGs that cannot be made planar and kept connected, simultaneously using standard *topology control* techniques. In our ADHOC-NOW 2007 [17] paper, we model the WSN as a *d-Quasi-unit disk graphs* (d-QUDG)[5] and address the problem highlighted above. We present a distributed topology control protocol that runs on a d-QUDG for $d \geq 1/\sqrt{2}$, and computes a sparse, constant-spanner, both in Euclidean distance and in hop distance. Our protocol is local and runs in $O(1)$ rounds of distributed computation. The output topology permits memoryless (geographic) routing with guaranteed delivery. In fact, when our topology control protocol is used as preprocessing step for the geographic routing protocol GOAFR$^+$, we get the routing time guarantee of $O(\ell^2)$ for any source-destination pair that are ℓ units away from each other in the input d-QUDG.

Scheduling to Minimize Communication Interference in Radio Networks

As noted earlier, due to the difficulty in performing regular maintenance on WSNs, battery reserve becomes a valuable resource after deployment. Hence, it is preferred to keep node activity to a minimum while accomplishing the tasks that nodes in the network are expected to perform as a collective. Due to the fact that primitive communication devices, such as small band radio, are used to maintain the networking aspect of WSN and the fact that in typical deployment scenarios the density of nodes with respect to their radius of communication might be quite high, cacophony results as nodes wake up. Hence, transmission is wasted in cases where receivers are not listening or where a receiver is the target of multiple simultaneous transmissions, leading to collisions. *Time Division Multiple Access* (TDMA) is a real-time scheduling technique to avoid such situations. Moreover, some popular sensor platforms consume as much power receiving as transmitting messages [22], so scheduling time-slots can also be used by receivers to further reduce energy needs by powering off the radio during idle periods. In WSNs, a mode of communication of interest is *unicast*. A unicast succeeds if the targeted neighbor does not concurrently transmit, and the target receives without collision. Radio transmissions in WSNs are not spatially directed. A unicast could be overheard by a number of transmitter's neighbors; and at a non-targeted neighbor, there can be a collision or that neighbor could be transmitting concurrently. The problem of scheduling unicasts is equivalent to assigning to each edge in the communication network a color so that transmissions on edges of the same color are concurrent and there are no collisions at

---

[5] A d-QUDG consists of vertices in $\mathbb{R}^2$ and an edge set E satisfying the rules: (i) $\{u,v\} \in E$ if $\|u-v\|_2 \leq d$ and (ii) $\{u,v\} \notin E$ if $\|u-v\|_2 > 1$. Note that edges between pairs of vertices u, v with $d < \|u-v\|_2 \leq 1$ are left unspecified and are assumed to be provided by an adversary.

unicast destinations. Gandham et al. [8] propose a distributed protocol to find such a link schedule for *acyclic networks* that runs in $\Theta(n)$ rounds and computes a link schedule using 2Δ time slots.

For large sensor networks, scheduling protocols with linear running time may be unacceptably slow. However, paying a small additional penalty in terms of time slots makes a $o(n)$ round scheduling algorithm more appealing. In our paper that appeared in SENSORWARE 2006 [12], we present a simple, randomized, distributed algorithm that runs in $O(\text{polylog}(n))$ rounds on acyclic networks and computes a link schedule using $2\Delta\cdot(1+\varepsilon)$ time slots, for any $\varepsilon > 0$.

k–Clustering to Minimize Sum of Radii

Given a metric space (V, d), define the ball $\mathrm{B}(v, r)$ centered at $v \in V$ and having radius $r \geq 0$ to be the set $\{q \in V | d(v, q) \leq r\}$. Given a finite metric space (V, d), and an integer $k > 0$, the *minimum-cost k-cover* aims to cover V with at most k balls centered at points in V while minimizing the sum of the radii of these balls. In the *metric* version of this problem, the metric space (V, d) is explicitly specified. In the Euclidean version, V is given as a set of points in some fixed dimensional Euclidean space $\mathbb{R}^l$, and d is then the standard Euclidean distance. Both the metric and the Euclidean version of the problem have been well examined, motivated by applications in clustering and base-station coverage [7, 5, 16, 3, 2].

In our SODA 2008 paper [10], we consider the Euclidean version of the problem in the plane. Assuming that two sums of square roots of rational numbers can be compared in polynomial time, we give an exact algorithm with polynomial running time. In a model where the two costs cannot be compared in polynomial time, we give a $(1+\varepsilon)$-approximation, for any $\varepsilon > 0$. The algorithm generalizes in a straightforward manner to any fixed dimension l and to some other related problems.

In our SWAT 2008 paper [11], we consider the problem in a more general setting – that of metric spaces. We give a randomized algorithm for the minimum cost k-cover for an n-point general metric space which runs in $n^{O(\log n \cdot \log \Delta)}$ (where Δ is the ratio between the largest and the smallest distance), and computes, with high probability, an exact solution. We also show that the dependence on Δ is necessary as we show that the problem is NP-hard on planar metrics and, more surprisingly, on metric spaces of constant doubling dimension.

Future Goals

My plan for the immediate future is to continue research in the field of design and analysis of distributed algorithms motivated by problems in wireless networks. I intend to expand on my current work by investigating the following problems: (i) devise exact algorithms for the *minimum-cost k-cover* on various classes of metric spaces; (ii) devise a polynomial time constant-factor approximation to the *domatic partition problem on UDGs*; and (iii) I would like to explore local solutions that give sharper bounds for the *UDG realization problem.*

While working at NDSSL/VBI[6] at Virginia Tech this summer, I implemented and tested a multicommodity maximum flow algorithm of Garg-Könemann [9] using the *Boost Graph Library* [1]. I would like to devise some natural heuristics from the algorithm to enable the algorithm to run on massive graphs as a test-bed for various simulations involving network flow. I would like to learn more about and be involved in designing algorithms for improving network throughput in large scale networks by only examining local network traffic patterns. This experience has given me a taste of experimental and implementation-based research related to wireless networks.

In the long run, I would like to continue to be involved in the design, analysis, and implementation of distributed algorithms for wireless networks. I would like to expand my work to models of wireless networks other than the graph-theoretic models that include specific engineering constraints and physical properties. I would also like to develop expertise in areas pertaining to streaming algorithms, bioinformatics, and large scale network simulations.

References

[1] *Boost Graph Library.* http://www.boost.org/libs/graph/.

[2] Helmut Alt, Esther M. Arkin, Hervé Brönnimann, Jeff Erickson, Sándor P. Fekete, Christian Knauer, Jonathan Lenchner, Joseph S. B. Mitchell, and Kim Whittlesey. Minimum-cost coverage of point sets by disks. In Nina Amenta and Otfried Cheong, editors, *Symposium on Computational Geometry*, pages 449–458. ACM, 2006.

---

[6] `http://ndssl.vbi.vt.edu`

[3] Vittorio Bilò, Ioannis Caragiannis, Christos Kaklamanis, and Panagiotis Kanellopoulos. Geometric clustering to minimize the sum of cluster sizes. In Gerth Stølting Brodal and Stefano Leonardi, editors, *ESA*, volume 3669 of *Lecture Notes in Computer Science*, pages 460–471. Springer, 2005.

[4] P. Bose, P. Morin, I. Stojmenovic, and J. Urrutia. Routing with guaranteed delivery in ad hoc wireless networks. *Wireless Networks*, 7(6):609–616, 2001.

[5] Moses Charikar and Rina Panigrahy. Clustering to minimize the sum of cluster diameters. *J. Comput. Syst. Sci.*, 68(2):417–441, 2004.

[6] B. Chen, K. Jamieson, H. Balakrishnan, and R. Morris. Span: An energy-efficient coordination algorithm for topology maintenance in ad hoc wireless networks. *ACM Wireless Networks Journal*, 8(5), 2002.

[7] Srinivas Doddi, Madhav V. Marathe, S. S. Ravi, David Scot Taylor, and Peter Widmayer. Approximation algorithms for clustering to minimize the sum of diameters. *Nord. J. Comput.*, 7(3):185–203, 2000.

[8] S. Gandham, Milind Dawande, and R. Prakash. Link scheduling in sensor networks: distributed edge coloring revisited. In *INFOCOM*, pages 2492–2501. IEEE, 2005.

[9] Naveen Garg and Jochen Könemann. Faster and simpler algorithms for multicommodity flow and other fractional packing problems. In *FOCS*, pages 300–309, 1998.

[10] Matt Gibson, Gaurav Kanade, Erik Krohn, Imran A. Pirwani, and Kasturi R. Varadarajan. On clustering to minimize the sum or radii. In *SODA*, 2008.

[11] Matt Gibson, Gaurav Kanade, Erik Krohn, Imran A. Pirwani, and Kasturi R. Varadarajan. On metric clustering to minimize the sum or radii. In *SWAT*, 2008.

[12] Ted Herman, Sriram V. Pemmaraju, and Imran A. Pirwani. Oriented edge colorings and link scheduling in sensor networks. In *In SENSORWARE 2006: 1st International Workshop on Software for Sensor Networks*, 2006.

[13] Brad Karp and H. T. Kung. Gpsr: greedy perimeter stateless routing for wireless networks. In *MobiCom '00: Proceedings of the 6th annual international conference on Mobile computing and networking*, pages 243–254, New York, NY, USA, 2000. ACM Press.

[14] Fabian Kuhn, Thomas Moscibroda, Regina O'Dell, Mirjam Wattenhofer, and Roger Wattenhofer. Virtual coordinates for ad hoc and sensor networks. To appear in *Algorithmica*, 2006.

[15] Fabian Kuhn, Roger Wattenhofer, Yan Zhang, and Aaron Zollinger. Geometric ad-hoc routing: of theory and practice. In *PODC '03: Proceedings of the twenty-second annual symposium on Principles of distributed computing*, pages 63–72, New York, NY, USA, 2003. ACM Press.

[16] Nissan Lev-Tov and David Peleg. Polynomial time approximation schemes for base station coverage with minimum total radii. *Computer Networks*, 47(4):489–501, 2005.

[17] Kevin M. Lillis, Sriram V. Pemmaraju, and Imran A. Pirwani. Topology control and geographic routing in realistic wireless networks. In Evangelos Kranakis and Jaroslav Opatrny, editors, *ADHOC-NOW*, volume 4686 of *Lecture Notes in Computer Science*, pages 15–31. Springer, 2007.

[18] T. Moscibroda and R. Wattenhofer. Maximizing the lifetime of dominating sets. In *Proceedings of the 5th. IEEE International Workshop on Algorithms for Wireless, Mobile, Ad Hoc and Sensor Networks*, 2005.

[19] Dragos Niculescu and Badri Nath. Ad hoc positioning system (APS). In *In Proceedings of Sixth Global Internet Symposium (GI2001) in conjunction with IEEE Globecom 2001*, pages 2926–2931, November 25-29 2001.

[20] Sriram V. Pemmaraju and Imran A. Pirwani. Energy conservation via domatic partitions. In Sergio Palazzo, Marco Conti, and Raghupathy Sivakumar, editors, *MobiHoc*, pages 143–154. ACM, 2006.

[21] Sriram V. Pemmaraju and Imran A. Pirwani. Good quality virtual realization of unit ball graphs. In Lars Arge, Michael Hoffmann, and Emo Welzl, editors, *ESA*, volume 4698 of *Lecture Notes in Computer Science*, pages 311–322. Springer, 2007.

[22] Joseph Polastre, Robert Szewczyk, and David Culler. Telos: enabling ultra-low power wireless research. In *IPSN '05: Proceedings of the 4th international symposium on Information processing in sensor networks*, page 48, Piscataway, NJ, USA, 2005. IEEE Press.

[23] Ananth Rao, Christos Papadimitriou, Scott Shenker, and Ion Stoica. Geographic routing without location information. In *MobiCom '03: Proceedings of the 9th annual international conference on Mobile computing and networking*, pages 96–108, New York, NY, USA, 2003. ACM Press.

[24] Santosh Vempala. Random projection: A new approach to VLSI layout. In *FOCS '98: Proceedings of the 39th Annual Symposium on Foundations of Computer Science*, pages 389–395, Washington, DC, USA, 1998. IEEE Computer Society.