R E S E A R C H   S T A T E M E N T   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. Much of my work has been on various geometric and combinatorial optimization problems in wireless networks. In the following, I describe my recent research by presenting 4 problems I have worked on lately. Following this, I describe a plan for the future[1].

Domatic Partition Problem: Some of my work has been driven by the need to extend the lifetime of wireless sensor networks (WSNs) by carefully scheduling node activity. Partitioning the nodes in a WSN into many disjoint *dominating sets* has been proposed as a model to extend the network life without sacrificing coverage responsibility. In our MobiHoc 2006 paper, we consider the problem of finding a large *domatic partition* on a *unit disk graph* (UDG) and present fast, distributed algorithms on a variety of network models. While we do not obtain a constant-factor approximation to the domatic partition problem on UDGs, we do obtain a vertex partition which could be turned into a contant-factor approximation if the radio signal at each node is boosted by a small additional amount.

Unit Disk Graph Realization: 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. In work that has appeared in ESA 2007, we give an algorithm that embeds a UDG without geometry in the Euclidean plane while minimizing the ratio of the longest edge to the shortest non-edge (also known as *quality* of *UDG realization*). While the quality of our solution is only marginally better than the best known solution (Vempala, FOCS 1998), the key feature of our solution is that it is completely combinatorial as opposed to solving exponentially large LPs using the *ellipsoid method* as previous solutions do. The solution depends upon exploiting certain key combinatorial properties of UDGs which have 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.

Guaranteed Geographic Routing Beyond UDGs: 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 paper, we model the WSN as a *d-Quasi-unit disk graphs* (d-QUDG)[2] 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.

Minimum Cost k-Clustering: Most recently I have worked on clustering problems also motivated by challenges in wireless networking such as *minimum cost k–clustering*. In our SODA 2008 paper, we consider the the problem in which we are given n points in the plane, and a positive integer k, and asked to find at most k disks so that the sum of the radii of all the disks is minimized and all n points are covered. 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 higher dimension.

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 besides the Euclidean space; (ii) devise a polynomial time constant-factor approximation to the *domatic partition problem on UDGs*; and (iii) explore local solutions that give sharper bounds for the *UDG realization problem*. In the long run, I would like 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 that include specific engineering constraints and physical properties. I would also like to develop expertise in areas pertaining to streaming algorithms, bio-informatics, and large scale network simulations.

[1]For a detailed version, please see http://www.cs.uiowa.edu/~pirwani/jobapplication/pirwani-research-statement.pdf

[2]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.