

## Dinesh P. Mehta: Research Statement

My research area is in applied algorithms and data structures. A typical paper consists of formulating an algorithmic problem, arguing that this formulation is practical in the context of the application area, developing theory to better characterize the problem, and then exploiting the theory to develop an efficient algorithm. When the problem is inherently intractable, as a lot of real-world problems are, pure algorithmic theory only takes you so far. In these cases, my research involves developing algorithms that do well in practice. Sometimes, these solutions find their basis, not in classical algorithmics, but in techniques such as simulated annealing and genetic algorithms or linear/integer programming that arise in other disciplines. Software that implements the algorithms is then developed and evaluated by running experiments on benchmarks. The performance of the algorithm is judged on runtime and on the quality of the solution it obtains. Thus, my research usually makes both a theoretical and an experimental contribution.

Some of the application areas I have studied relatively recently or am currently studying are listed below:

1. *VLSI Floorplanning*. Floorplanning is the placement of large circuit components on a chip to optimize various criteria and is an integral step in the design of modern chips. Most of my more recent work in this area is joint with Yan Feng, who received his PhD in 2004. Yan's research resulted in two publications in IEEE Transactions on CAD and was presented at three conferences. Our contribution was to formulate the floorplanning problem (which has been around for several decades) in a manner that is more suitable to industry (our formulation was developed in collaboration with scientists at Intel) and then solve our version of the problem using a combination of graph-theoretic, geometric and iterative techniques. Our algorithms were implemented and run on established benchmarks. Yan is now at Cadence after a stint at U. Minnesota as a post-doctoral researcher with Prof Sapatnekar.
2. *Mobile computing*. My work in this area is with Sahar Idwan, who received her PhD in 2005. Sahar's work is on the pursuer-evader problem. This problem enjoyed a resurgence because it was posed as a challenge problem by the DARPA NEST (Network Embedded Software Technology) program. The objective of this research was for the pursuer to capture the evader in minimum time using a minimum number of updates in a road network. Sahar's research involved developing lower bounds and approximation algorithms for the problem, developing probabilistic graph-based algorithms and employing hierarchical graphs to improve computation time. Algorithms were developed using LEDA and were run on several benchmarks including a roadmap of the Denver metropolitan area. This resulted in one journal and three conference publications. Sahar is currently an Assistant Professor at Hashemite University in Jordan.
3. *Cheminformatics*. My earlier work in this area is with John Crabtree, who received his PhD in 2007 and with Professors Dean and McKinnon in Chemical Engineering. My more recent work is with Tina Kouri, a PhD candidate in our department. The overall

goal of this research is to automate the analysis of large mechanisms that describe the operation of industrial chemical processes. At the heart of this is Automated Reaction Mapping (ARM), a difficult problem rooted in graph-isomorphism (note that graphs arise in this context because they can be used to model molecular structure). In his dissertation, John developed algorithms for ARM that were published in the ACM Journal of Experimental Algorithms. The algorithms were incorporated into a proof-of-concept software system that also includes a rule-based component to classify chemical reactions. A description of the system appeared in the Journal of Chemical Information and Modeling, a journal of the American Chemical Society. For her dissertation, Tina is developing order-of-magnitude faster solutions for ARM (making our algorithms suitable for larger bio-mechanisms) and is applying, refining, and “hardening” our classifications of reactions to well-established mechanisms based on user feedback from chemical engineers. Our goal is to ultimately create a tool that chemical engineers actually use.

4. *Computational Materials.* My work in this area is with Keith Hellman (a PhD candidate in our department), whom I am co-advising with Professor Ciobanu in Engineering. Professor Ciobanu has successfully used genetic algorithms in the past to obtain a number of optimal nano-structure configurations. We are seeking to expand the scope and applicability of his optimization techniques by dramatically improving the run-time of the GA code - Keith plans to use a combination of improved system design, parallel computing, and machine learning to enable the GA algorithm to converge faster.
5. *Cyber-Enabled Efficient Management of Structures.* I am a member of an interdisciplinary team that also includes Electrical and Mechanical Engineers that was recently awarded a \$1.4M NSF Cyber-Physical Systems (CPS) grant. The purpose is to use and combine a number of techniques arising in different disciplines to optimize energy usage in buildings given environmental conditions and human activity patterns. This work is still in the early stages. I am currently working on formulating optimization problems that seek to obtain a desired temperature profile in the rooms of a building with minimum energy consumption. I am also part of a group with Profs. Hoff and Szymczak that is trying to deduce human activity patterns in a building from sensor data.