D I M A C S

Center for Discrete Mathematics and Theoretical Computer Science

DIMACS SPECIAL YEAR ON NETWORKS: 1997 ANNUAL REPORT

DIMACS Center • Rutgers University • 96 Frelinghuysen Road • Piscataway, NJ 08854-8018 • USA
TEL: 732-445-5928 • FAX: 732-445-5932 • EMAIL: `center@dimacs.rutgers.edu`
WWW: `http://dimacs.rutgers.edu/` • TELNET: `telnet info.rutgers.edu 90`

DIMACS is a partnership of Rutgers University, Princeton University, AT&T Labs - Research, Bell Labs, NEC Research, and Telcordia Technologies (formerly Bellcore)

Overview:

The Special Year on Networks began in July 1996 and will continue through August 2000 with the help of a three year grant received from the National Security Agency. The program is stressing postdoctoral training and workshops that will reach a broad segment of the DM and TCS research community as well as engage practitioners.

Organizers:

- Stuart Haber, Bellcore / Surety Technologies
- David Johnson, AT&T Labs - Research
- Mihalis Yannakakis, Bell Labs

Scientific Advisory Board: Baruch Awerbuch, Johns Hopkins; Br. Badrinath, Rutgers; Endre Boros, Rutgers; Ding-Zhu Du, Minnesota; David Goodman, Rutgers; Jack Holtzman, Rutgers; Frank Hsu, Fordham; Srinivasan Keshav, Cornell; Arjen Lenstra, Citibank; Udi Manber, Arizona; Ray Miller, Maryland; Debasis Mitra, Bell Labs; Clyde Monma, Bellcore; Panos Pardalos, Florida; K.K. Ramakrishnan, AT&T Labs - Research; Greg Shannon, Spanning Tree Technologies; Scott Shenker, Xerox; Subhash Suri, Washington U. in St. Louis; Eva Tardos, Cornell; Peter Winkler, Bell Labs; and Yacov Yacobi, Microsoft

Sponsors:

The National Science Foundation
The New Jersey Commission on Science and Technology
The National Security Agency
AT&T Labs - Research
Bell Labs
Bellcore
Rutgers University
Princeton University

Theme of program:

This special year is intended to bring together researchers and practitioners in a wide variety of fields connected to the common theme of "Networks" and ranging from questions of infrastructure to ones of applications. Special emphasis is placed on Internet-related topics, as this is a young and growing area where theoretical computer science and discrete mathematics have a chance to make a real impact.

The special year's topics can be grouped into four overlapping areas, which we list below along with examples of planned or possible workshop topics.

- **Network Security:** protocols for electronic commerce, public key infrastructure, network security and threat models.
- **Network Applications:** information retrieval, databases and the Web, multimedia transmission, and coding distributed and collaborative computing.
- **Network Control:** routing, congestion and admission control, network protocols, mobile computing.
- **Network Design:** topology and connectivity, switch design, facilities location.

The goal is to engage theoreticians with these problems and foster communication between them and real-world practitioners. Practitioners have much to tell theorists about what the key issues are, and theorists may themselves be able to provide crucial insights and solutions to practitioners.

Workshops, Miniworkshops and Special Events 1997:

- **Workshop on Architecture and Algorithmic Aspects of Communication Networks**
  - Dates: January 26-28, 1997
  - Location: Rutgers University
  - Attendance: 70
  - Organizers: Eva Tardos, Cornell University; Serge Plotkin, Stanford University; Israel Cidon, Technion and Sun Microsystems

  The size of the future networks, combined with the need for bandwidth reservation and QoS guarantees, poses numerous challenging problems. These problems are exacerbated by the convergence of the data, telephone and video services that have to be supported in the same framework. In parallel, the recent advances in theoretical computer science and discrete mathematics bring us closer to being able to address many of these issues. Unfortunately, these advances have not had a significant impact on practical networks, mostly because of a lack of continuous dialog between the research communities.

  The aim of the workshop was to bring together and allow the exchange of research ideas and interests between experts in networking, theory, ATM and Iternet. Both theoretical and applied issues were considered. Primary areas of focus of the workshop included:

  - Switching: Packet scheduling and QoS guarantees, switching architecture.
  - Routing: Link state algorithms, dynamic online routing of virtual circuits, generation and maintenance of routing tables, rerouting due to bandwidth fragmentation, service restoration after link failure.
  - Network topology design.
  - Multicast: Membership management, multicast routing, QoS in multicast, policing.
  - Admission Control.
  - Fairness.

  Most of the workshop was devoted to talks by distinguished invited speakers but there were also a few contributed talks and an hour long panel discussion.

- **Workshop on Economics, Game Theory, and the Internet**
  - Dates: April 18-19 1997
  - Location: Rutgers University
  - Attendance: 80
  - Organizers: Eric Friedman, Rutgers; Scott Shenker, Xerox PARC

  The Internet poses many interesting challenges for theoretical economics and game theory. The highly asynchronous interactions and large degree of anonymity inherent in the Internet require that many traditional game theoretic notions of equilibrium, learning, and mechanism design need to be re-examined in this context. New cost sharing issues (such as when receivers share multicast transmissions, or when flows share the use of a link) arise in network settings. Incentive issues must be considered when designing basic network algorithms and protocols. In addition, regulation of the Internet and other telecommunication networks raises other theoretical issues. This workshop focused on identifying new areas of theoretical interest that arise out of the institutional and technical nature of the Internet and other communication networks.

Each day of the two day workshop began with two 90 minute plenary talks followed by more standard 30 minute research reports in the afternoons. Plenary speakers included Herve Moulin (Duke University) who discussed the applications of fair division, which has been much studied in economics, to problems in the Internet; Bill Sharkey (Federal Communications Commission) who described the Telecommunications Act of 1996 and its implications; Hal Varian (University of California - Berkeley) on intellectual property rights in a digital environment; and Scott Shenker (Xerox PARC) who gave a tutorial on the structure of the Internet and a review of important open problems.

The speakers were quite diverse, coming from Computer Science, Economics, Engineering, and Physics; they covered topics such as economic analyses of the Internet, game theoretic learning on the Internet, science policy related to the Internet, and general analyses of networks.

- **Workshop on Network Design: Connectivity and Facilities Location**
  - Dates: April 28-30, 1997
  - Location: Nassau Inn, Princeton, NJ
  - Attendance: 70
  - Organizers: Ding-Zhu Du, University of Minnesota; Panos Pardalos, University of Florida

  Connectivity and facilities location are two important topics in network design with applications in data communication, transportation, production planning, and VLSI designs. There are two issues concerning these two topics: design and optimization. They involve combinatorial design and combinatorial optimization. Finding the solution of design problems and the optimal or approximate solution of the related optimization problem are challenging tasks because no polynomial time algorithms are known. Such problems include some variations of Steiner tree problems (such as multiple connected Steiner network, independent flow problem, and subset-interconnection designs), topology network design, nonlinear assignment problems (such as quadratic assignment problems), problems in facilities location and allocation and network problems appearing in VLSI design. The workshop focused on combinatorial, algorithmic, and applicational aspects of these problems.

  A refereed volume with papers presented at the Workshop will be published early next year by the American Mathematical Society in the DIMACS series.

- **Workshop on Network Switching**
  - Dates: July 7-9, 1997
  - Location: Princeton University
  - Attendance: 17
  - Organizers: Frank Hwang, AT&T; Ding-Zhu Du, University of Minnesota

  Recent development in Internet has required new technology in network switching. This includes routing algorithm designs and network designs to meet various communication requirements. This workshop focused on all aspects of network switching, including the following topics:

- Network Environment: circuit-switching, packet switching, ATM switching, multirate switching, wireless.
- Routing: point-to-point, permutation, broadcast, conferencing.
- Network Topology: multistage, rings, cubes, grids, stars, general connectors, fault-tolerant networks.
- Switching Component: concentrator, superconcentrator, partial concentrator, expander, multiplier, copy network, n-generalizers, optical network.
- Nonblockingness: strictly nonblocking, wide-sense nonblocking, rearrangeable.
- Optimization: hardware minimization, optimal network design.

A refereed volume with papers presented at the Workshop will be published the next year by the American Mathematical Society in the DIMACS series and selected papers will also be published in a special issue of *Networks.*

- **Design and Formal Verification of Security Protocols**
  - Dates: September 3-5, 1997
  - Location: Rutgers University
  - Attendance: 110
  - Organizers: Hilarie Orman, DARPA; Catherine Meadows, Naval Research Laboratory

Research has been intensifying in the application of formal methods to cryptographic protocol verification. The goal of this workshop was to facilitate this process by bringing together those who are involved in the design and standardization of cryptographic protocols, and those who are developing and using formal methods techniques for the verification of such protocols.

The workshop was organized as a series of related talks by individuals and panels of experts, scheduled with ample time for questions. The goal was to encourage audience participation in order to foster understanding of the technology and how it could be applied, and this supported the overall goal of getting researchers and protocol designers to exchange ideas on how to make substantial progress in this field. We alternated papers with panels soliciting new paths for research with emphasis on proposals addressing new protocols with respect to their formal and informal analysis.

Other topics of interest included:

- Progress in belief logics
- Use of theorem provers and model checkers in verifying crypto protocols
- Interaction between protocols and cryptographic modes of operation
- Methods for unifying documentation and formal, verifiable specification
- Methods for incorporating formal methods into crypto protocol design
- Verification of cryptographic API systems
- Formal definition of correctness of a cryptographic protocol
- Arithmetic capability required for proofs of security for number theoretic systems
- Formal definitions of cryptographic protocol requirements
- Design methodologies
- Emerging needs and new uses for cryptographic protocols

  - Multiparty protocols, in particular design and verification methods

  The workshop papers were published on-line the day before the workshop started, and a CD-ROM with the papers was produced as well.

- **End-to-End Network Modelling and Simulations**
  - Dates: October 23-24, 1997
  - Locations: Rutgers University
  - Attendance: 100
  - Organizers: Sandeep Bhatt, Bellcore; Andrew Ogielski, Rutgers; Ashok Erramilli, QMetrix

  This workshop focused on current developments in modeling, simulation, and algorithmic techniques which take into account the observed and measured behavior of networks. In order to obtain a reasonably complete picture of the state of the art, and to stimulate further research, the organizers invited practitioners of industrial and defense network modeling together with researchers involved in the analysis and simulations of networks and protocols. They were particularly interested in identifying areas requiring further research, and in creating an informal forum for practitioners, researchers, and students to hear each other and to exchange ideas.

  The primary topics covered in the workshop were:

  - Network modeling and simulation practices
  - Modeling and simulation of wireless networks
  - Models of traffic sources and aggregated traffic
  - Modeling and simulation of network protocols
  - Parallel simulation techniques for large telecommunication systems
  - Algorithmic analyses of network and protocol models

  The demographic mix of the participants (20% were from corporations unaffiliated with DIMACS, 40% were from corporations affiliated with DIMACS (AT&T, Bellcore, Lucent) and 40% were from academia and national laboratories) indicates that we have succeeded in creating an atmosphere conducive to university-industry collaborations, and in reaching out to the networking and communications industries.

  All workshop viewgraphs will be published on the DIMACS web pages together with extra papers and links provided by the speakers.

- **Networks in Distributed Computing**
  - Dates: October 27-29, 1997
  - Location: Rutgers University
  - Attendance: 50
  - Organizers: Marios Mavronicolas, University of Cyprus; Michael Merritt, AT&T Labs - Research; Nir Shavit, Tel Aviv University and MIT

  The primary goal of the Workshop was to bring together researchers from both industry and academia interested in issues related to networks as they appear in the context of distributed computing. We invested special effort to expose the theoretical community carrying out basic research in the field of distributed computing to fundamental problems and challenges related to recent technological advances in the networking industry. We invited both theoreticians and practitioners, in a hope that the latter would carry through their practical motivation, experience, and insights to the Workshop

Topics of discussion included: ATM networking technology, issues of routing and flow control in modern communication networks, service pricing, security, optical networking, management of Internet, mobile computing, and formal verification of communication protocols.

A Selection of Special Year Research Results

The following **abbreviations** are used throughout: RU = Rutgers University, PU = Princeton University, ATTL = AT&T Labs, BL = Bell Labs, BC = Bellcore, NEC = NEC Research Institute, PM = permanent member of DIMACS, PD = DIMACS postdoctoral fellow, GS = Graduate Student, V = DIMACS visitor.

A Local Search Heuristic for Facility Location Problems Madhukar Korupolu, Greg Plaxton, and Rajmohan Rajaraman (PD) have considered approximation algorithms for several NP-hard problems arising in facility location. In each problem, we are given: (i) a set of locations at which we may open a facility, (ii) a set of client locations, together with their demands, that require to be serviced by a facility, and (iii) for every pair of locations the cost of shipping a unit of demand between the two locations. In the *k-median problem*, the objective is to determine a set of at most k locations at which to open a facility so as to minimize the total shipping cost. In the *facility location problem*, we are also given a cost of opening a facility at any location, and the objective is to determine a set of locations at which to open facilities so as to minimize the total shipping and opening costs. We differentiate between the uncapacitated and the capacitated variants of the above problems. In the uncapacitated case, an unlimited amount of demand can be shipped from any facility, whereas in the capacitated case, there is an upper bound on the total demand that can be shipped from any facility. Koropolu, Plaxton, and Rajaraman proved that a simple local search heuristic yields polynomial-time constant-factor approximation bounds for the metric versions of the uncapacitated k-median and facility location problems. (For the k-median problem, their algorithms require a constant-factor blowup in the parameter k.) This local search heuristic was first proposed several decades ago, and has been shown to exhibit good practical performance in empirical studies. Koropolu, Plaxton and Rajaraman also extended the above results to obtain constant-factor approximation bounds for the metric versions of the capacitated k-median and facility location problems.

A Simple Distributed Algorithm for Placing Objects in Hierarchical Networks Madhukar Korupolu, Greg Plaxton, and Rajmohan Rajaraman (PD) have designed a simple algorithm to determine a placement of objects in hierarchical networks. Their placement has the property that the expected access time of a request is within a small constant factor of the optimal. Moreover, the placement is obtained by a single bottom-up pass of the hierarchical network and hence can be implemented in a distributed manner. The algorithm has practical potential in the design of World Wide Web caching schemes and in allocating files in distributed filesystems.

Nearly Linear Time Approximation Schemes for Euclidean TSP, Steiner Tree and other Geometric Problems Much effort has been invested in designing approximation algorithms for NP-hard geometric problems such as Euclidean Traveling Salesman and Minimum Steiner Tree. In 1996, Sanjeev Arora (PM, PU) presented new algorithms for many of these problems that are approximation schemes: for any fixed $c > 0$, the algorithm computes a solution of cost at most $1 + c$ times the optimal. The running time was $n^{O(1/c)}$ in the plane and slightly superpolynomial in higher dimensions. (Mitchell independently discovered a similar scheme for the planar case.) Arora has now found a new approximation scheme that is simpler and much more efficient. For any fixed $c > 0$, it finds a $(1+c)$-approximate solution in $n * \text{polylog}(n)$ time, which is nearly linear in n. The scheme also works in nearly linear time on instances in d dimensions (for fixed d). His algorithm generalizes to many other geometric problems, including the inimum Steiner Tree problem, and the Euclidean Matching problem.

A New Model of Generalized Steiner Networks J.F. Weng (V) studied 2-dimensional networks with a Steiner topology. Suppose T is such a network. If any edge e joining two Steiner points s and s' satisfies the condition that the other edges of s, s' are, in pairs, parallelly lying on two sides of e, then T is called a *generalized Steiner network*. As an example, Weng showed that the drive shafts for an underground mine form a network of this type since the shafts have to have a maximum gradient. He gave an algorithm that can compute generalized Steiner networks with given Steiner topologies in linear time.

Compliance Checking in the PolicyMaker Trust Management System Emerging electronic commerce services that use public-key cryptography on a mass-market scale require sophisticated mechanisms for managing trust. For example, any service that receives a signed request for action is forced to answer the central question "Is the key used to sign this request authorized to take this action?" In some services, this question reduces to "Does this key belong to this person?" In others, the authorization question is more complicated, and resolving it requires techniques for formulating security policies and security credentials, determining whether particular sets of credentials satisfy the relevant policies, and deferring trust to third parties. Blaze, Feigenbaum, and Lacy identified this "trust management problem" as a distinct and important component of network services and described a general tool for addressing it, the "PolicyMaker trust management system." At the heart of a trust management system is an algorithm for "compliance checking." The inputs to the compliance checker are a request, a policy, and a set of credentials. The compliance checker returns yes or no, depending on whether the credentials constitute a "proof" that the request complies with the policy. Thus a central challenge in trust management is to find an appropriate notion of proof and an efficient algorithm for checking proofs of compliance. Stimulated by discussions at the DIMACS workshop on Trust Management in Networks, Matt Blaze (PM, ATTL), Joan Feigenbaum (PM, ATTL), and Martin Strauss (PM, ATTL) have developed the notion of proof that is used in the current version of the PolicyMaker trust management system. They showed that this notion of proof leads to a compliance-checking problem that is undecidable in its most general form and is NP-hard even if restricted in several natural ways. They identified a special case of the problem that is solvable in polynomial time and is widely applicable. The algorithm that they gave for this special case has been implemented and is used in the current version of the PolicyMaker system.

Reliable Communication over Partially Authenticated Networks Reliable communication between parties in a network is a basic requirement for executing any protocol. Amos Beimel (PD) and Matt Franklin (PM, ATTL) considered the effect on reliable communication when some pairs of parties have common authentication keys. The pairs sharing keys define a natural "authentication graph", which may be quite different from the "communication graph" of the network. They characterized when reliable communication is possible in terms of these two graphs, focusing on the very strong setting of a Byzantine adversary with unlimited computational resources. More specifically, reliable communication involves a transmitter which sends a message to a receiver. There are two versions of this problem: "single-pair" and "all-pairs". For single-pair reliable communication, one specifies the transmitter and receiver, and wants to succeed against any coalition of at most t

faulty processors. For all-pairs reliable communication, one wants to succeed for any transmitter, any receiver, and any coalition of at most t faulty processors. The network of channels defines a natural "communication graph", with an edge between two vertices for every channel between two processors. The pairs of parties sharing authentication keys define a natural "authentication graph", with an edge between two vertices for every shared key. Success for the all-pairs reliable communication problem depends in a natural way on the connectivity of these graphs. Beimel and Franklin showed that all-pairs reliable communication is possible if and only if the communication graph is $(t+1)$-connected and the union of the two graphs is $(2t+1)$-connected. When there is a specific sender a and a specific receiver b, the situation is not so straightforward. In this case the characterization involves a family of inductively defined graphs, each one containing all edges of the communication graph and some of the edges of the authentication graph.

Escrowed Identity Joe Kilian (PM, NEC), and Erez Petrank (PD) worked on secure authorized access. The issue that arises is the tension between the need of a person to access services without revealing his identity, and the need to reveal identities of users in extreme cases, for example, when a crime is committed. Kilian and Petrank used Key escrow techniques to meet both requirements. Key escrow has proven an active and contentious field of research and discussion. Essentially all of the attention in this area has been restricted to the simple case of communication: party A sends an encrypted message $E_K(M)$ to party B; some centralized authority is given the capability to recover either K or the specific message M. Kilian and Petrank considered a new domain for the application of key escrow ideas: Escrow in identity schemes. One distinctive feature of this application is that Escrowed identity may actually enhance privacy. By default, many identity schemes often require a person to give their entire identity "up front." A protocol in which this information is only released under special circumstances may prove an acceptable, and more private, substitute. Another distinctive feature is that Escrowed identity schemes work to the advantage of at least one of the parties invoking them. In traditional key escrow systems, both parties only lose by following the escrow system, and have everything to gain by bypassing it (which is generally quite easy to do). The motivation for traditional identification schemes in access control, communication and commerce need not be discussed. A rapidly increasing number of interactions take place on the purely information level, and are likely to become ubiquitous. As just one example, on some toll bridges one can now electronically identify oneself to a toll booth instead of waiting in line to pay a toll. Although this application is cost effective (or will be) and more convenient, it raises critical issues of privacy; a more pervasive toll collection system might as a side effect allow the tracing of people's movements to an unprecedented degree. Kilian and Petrank addressed the tension between party A's desire for privacy and anonymity and B's (and society's) desire to know who they are interacting with. This balance point of this tension can shift under certain often rare circumstances. For example, consider an automated parking garage. What is paramount to the garage is that the person entering it is authorized to do so; the person's precise identity is normally not a valid concern. However, suppose that on some night a person was murdered in the garage. At this point, the garage owner and society at large may have a legitimate interest in knowing who was there on that night. Or, one might wish to have a computer "chat room" in which one has conditional anonymity messages: As long as one follows the rules laws, one's identity is secure from even the system administrator. But if one flagrantly breaks the rules (such as arranging drug deals in the "Lion King" kiddie chat room), suitable law enforcement agencies can be appealed to in order to determine one's identity. Traditionally, access control has been all or nothing. One obtains all the information about the other person up front, with no recourse to learn more if circumstances warrant. This rigidity generally leads one to allow less privacy, since one is likely to want as much information as one can get just in case a "bad case" arises. Kilian and Petrank have given a more flexible alternative. They considered a more flexible, two-tier approach to identification. On the lower tier, a person gives only as much information about themselves as is strictly necessary for ordinary circumstances. In the garage example above, this means proving authorization to use the garage. On the second tier, a person gives a more precise statement of their identity than may be needed in extraordinary circumstances. This second tier is only accessible with the help of a third party, i.e., the escrow agency, which is preferably separate from and not under the control of the party managing access (such as the garage owners). An escrowed identity system consists of the following parties: An Identifier (the party who identifies themselves to the verifier, the user of the garage in the above example); An Issuer (the party which issues certificates that allow the identifier to identify themselves in an escrowed manner); A Verifier (typically the access provider who verifies the first-tier identity process as well as the escrow proof for the second-tier identity (the garage keeper in the above example); and Escrow Agent(s) (parties who use information forwarded by the verifier to make a second-tier identification of the identifier. These parties execute the following protocols: Issuing a certificate; checking the weak identity; recovering the complete identity. Kilian and Petrank have developed the properties of identity escrow systems. A typical requirement is that faking a legitimate identity is hard: A computationally bounded user that has not been issued a certificate by the issuer will fail to convince the verifier of having a certificate with probability almost 1. A second requirement is that ducking the escrow is hard: The probability that a computationally bounded user passes the check for weak identity but his identity cannot be recovered from the transcript by the escrow agents is almost 0. Kilian and Petrank have developed two concrete implementations of escrow identity systems. One is based on the RSA encryption scheme and the other on the El-Gamal signature and encryption schemes. An international patent is pending on this work (by Rutgers University and NEC).

Message Authentication Message authentication is one of the most used cryptographic primitives. In this setting, two parties which share a secret key would like to make sure that the messages they exchange are not changed on the communication channel before delivery. The Cipher Block Channing (CBC) Message Authentication Code (MAC) is an authentication method which is widely used in practice. Nevertheless, until recently no solid theoretical approach was suggested for proving security of such systems. The first paper to provide rigorous proof of security for CBC MAC (conditioned on the security of the underlying cryptographic block cipher) was by Bellare, Kilian and Rogaway in 1994. They suggested a framework for proving security of CBC MAC, and also used this framework to prove the security of CBC MAC follows from the security of the underlying block cipher when the messages have fixed length. They also suggested variants of CBC MAC that handle variable length messages but the length of the message has to be known in advance. Erez Petrank (PD) and Charles Rackoff followed their work by studying the case in which the length of the message is not known until the message ends. This seems an important issue in particular because the transmission of fax is inherently like that: When the transmission starts, the machine does not know how many pages are going to be input. Petrank and Rackoff first presented a variant of CBC MAC, called *double MAC* (DMAC), which handles variable unknown lengths. They proved that its security is implied by the security of the underlying block cipher. Next, they argued that the basic CBC MAC is secure when applied to prefix free message space. A message space can be made prefix free by authenticating also the (usually hidden) last character which marks the end of the message. This character never appears inside the message.

Associated Projects on Network Modeling and Parallel Simulations

In addition to the research work arising from its Networks special year, DIMACS has has two multi-disciplinary, multi-institutional research projects in network modeling and parallel simulations, each funded separately, but closely related to our Special Year in Networks. These projects relate to some very important problems in networking research, aiming to provide the methodology and software tools for creating software descriptions of very large networks – such as the global Internet – which could then be simulated on multiprocessor machines or on multiple networked machines, to speed up the simulations. Parallel discrete event simulations of telecommunications systems involve very difficult problems, and so far have not been successfully implemented in commercial systems despite a great desire for such tools from industry and the military.

Project on Simulations of Integrated Communications Systems

Project Summary: The centerpiece of this project is the design, prototype implementation and extensive testing of a standard description methodology for various telecommunication systems, with the following properties: 1) Efficiency and flexibility; 2) Suitability for (semi)automatic scalable parallel execution of simulations; 3) Portability and independence of the parallel simulation platform. To achieve these goals, the project involves research on parallel discrete event simulation algorithms and software, and on network modeling. The effort on modeling feeds back new ideas and corrections to software design.

The projected impact of this project is best described by analogy with the extremely successful use of computer–aided design tools in the integrated circuit (VLSI) industry, where description methods combining design with simulations (expressed in description languages such as VHDL, SDL or Verilog) have allowed designs of ever higher levels of complexity. In contrast, network protocol design and engineering – which deals with systems of larger complexity than VLSI – has not yet fully adopted such tools, and not for large scale designs, primarily because simulations are: 1) Much more difficult (involving modeling of network user behavior and traffic in addition to multiple layers of protocols), and 2) Much more computationally demanding (the stochastic nature of the problems coupled with the large size of interesting systems often leads to intractably long simulations).

The observed inefficiencies and recurring failures of very large data networks such as the global Internet show the desirability of such tools. This project is funded by the National Science Foundation for a three year period over 1995 - 1998, for a total of $1.5M.

Participating Scientists: Sandeep Bhatt (PM, BC and RU), Richard Fujimoto (Georgia Tech), Albert Greenberg (PM, ATTL), Jim Kurose (University of Massachusetts), David Nicol (William and Mary / Dartmouth), Andrew Ogielski (PM, RU), Diane Souvaine (PM, RU), Don Towsley (University of Massachusetts), and Paul Wright (PM, ATTL).

Collaborations: This project brings together experts in network modeling, design and analysis, distributed algorithms, and high performance discrete-event simulation to focus on the development of high-performance simulation tools for broadband networks and wireless communication services. This project directly brings together the expertise of scientists from two of the nation's leading telecommunications companies, in partnership with leading academic sites; and includes graduate and postdoctoral training that benefits students through exposure to industrial scale problems and internships.

Project Description: Early in the project the investigators developed a prototype network-centric description framework expressed as a a meta-language extending C++, called TeD (Telecommunications Description Language). TeD has parsimoniously incorporated the core modeling concepts identified over the years of research on hardware description languages, while at the same time providing facilities for parallel simulations without involving the modelers in the internal complexities of parallel simulators. Recognizing that modeling language design is an iterative process requiring user feedback, TeD and a compiler to translate TeD descriptions to executable parallel simulation code have been delivered to model developers in continually improving six releases (from 1.0 to the latest 1.6). Curently TeD compiles to the Georgia Time Warp (GTW) parallel simulation engine. The parallel simulations are performed on multiprocessor SGI and UltraSparc machines at Bellcore, Georgia Tech and University of Massachussets.

The investigation of the TeD methodology involves complex network modeling and parallel simulations subprojects: routing in broadband ATM PNNI internetworks; reliable multicast services; multipath propagation modeling in wireless networks.

The Private Network-Network Interface (PNNI) protocol suite is an international draft standard proposed by the ATM Forum. It is a complex set of ATM routing and signaling mechanisms and represents one of the most sophisticated signaling protocols devised to date, aimed at supporting QoS-based routing claiming unprecedented levels of scalability. This project created implementations of PNNI protocols and network configurations, with nodes stochastically generating requests for connection (VC) setup and teardown to test the actual performance of the protocols. Simulations were investigated on sample network configurations representative of regional telephony backbone networks, with hundreds of complex nodes, and demonstrated an unparalleled capability of simulating in parallel over one million of VC setups per wallclock hour.

Multicast is recognized as a fundamental operation in evolving Internet and in future networks. We identified the parallelism in modeling of multicast protocols and ran simulation experiments that demonstrate that. The available parallelism is a function of the degree to which the broadcast trees for multicasts overlap spatially and temporally. We have used TeD to evaluate two different approaches to providing reliable, scalable multicast communication: The *sender-initiated* and *receiver-initiated* approach. We have successfully built TeD performance models for both a specific sender-initiated protocol and two receiver-initiated protocols. We have successfully validated the simulations predicted throughput values against those predicted by analysis for up to 500 receivers.

Our efforts in the wireless area have focused on the design of modeling software that can provide radio propagation prediction at various levels of detail. This work is being undertaken collaboratively with researchers at WINLAB at Rutgers University. The project investigated parallel implementations of radio propagation models and interference calculations, and developed tools for ray-tracing calculation of multipath radio propagation in specified geographical environments, improving various computational geometry algorithms in the process.

As a complement to this project, we organized the DIMACS Workshop on End-to-End Network Modeling and Simulations at Princeton (Oct 23-24, 1997). The Workshop brought together the leading researchers on modeling the long–range dependent traffic and user behavior, on Internet modeling, on network simulations, and wireless network modeling.

Project on Scalable Self-Organizing Simulations (S3)

Project Summary: This is a project to build and distribute to the networking research community and the DoD a radically new, highly efficient network modeling and parallel simulation software package, building on the research carried out in the NSF project described in Section 5.8.1, and extending it in several important directions. The goal of this project is to achieve a generational transition in performance and manageability of simulations of telecommunication systems. This project is funded by DARPA for a three year period over 1996-1999, for a total of $3.1M.

Participating Scientists: Sandeep Bhatt (PM, BC and RU), Marina Chen and Mark Crovella (Boston University), James Cowie (Cooperating Systems Corporation), Richard Fujimoto, Kalyan Perumalla and Ellen Zegura (Georgia Institute of Technology), Don Leskiw (Ultra Corporation), David Nicol (Dartmouth), Andrew Ogielski (PM, RU).

Collaborations: The daunting complexity of modern communications technologies on the one hand, and of the large software systems required on the other, demands that a project of this sort be carried out by an interdisciplinary team of experts in parallel simulations, networking, software engineering, and defense applications. We have put together a team of scientists and engineers from six institutions to work on the project, including participants from two small companies. Academic collaborators involve in the project postdoctoral fellows and many graduate research assistants.

Project Description: Project S3 aims to provide the fastest parallel simulations, convenient modeling tools, and databases of scalable models for a range of networking problems, encompassing very different types of simulations such as wired network traffic, complex protocols stacks, packet interference in mobile wireless networks and military communications scenarios.

These goals require the common modeling and runtime environment (SSF, the **S**calable **S**elf–organizing **F**ramework) that allows *dynamic* hierarchical modeling suitable for parallel simulations, and allows one to simulate software models – without changes – both on conservatively and optimistically synchronized, radically new parallel simulators. The SSF API is the core interface common to all the individual S3 subprojects. Further speedup of simulations is achieved by researching and implementing runtime aggregation and deaggregation techniques, consisting of automated substitution of coarser (aggregate) submodels where absolute fidelity is not required.

The S3 project develops a truly process oriented, conservatively synchronized simulation engine **Nops**, the **No**rthern **p**arallel **s**imulator. Nops uses a modification of the threaded parallel runtime system, Cilk, developed at MIT. Cilk automates load-balancing through a process migration mechanism called "stealing". The first tests of Nops show excellent performance, providing a very strong proof of concept for this novel architecture of a parallel simulator.

The second parallel simulator, **MegEnta** (the name refers to a Million Entity Framework) consists of two distinct modules: TiCOON – A high-performance parallel simulation executive based on a novel micro-kernel approach which incorporates a Time Control mechanism that can simultaneously support different synchronization schemes; and HOPPERS – A novel ultra lightweight process abstraction which supports shared-memory semantics. Multiple instances of Nops and MegEnta parallel simulators can be linked over the network to create distributed federations of simulators.

Usability of parallel network simulators is a major goal of the S3 Project. Usability emerges from the separation of modeling work from the internal complexities of parallel simulators, and from the availability of user–friendly tools: for model browsing and assembly (CAD), for model debugging, for parallel performance monitoring and load management. The modeling environment and tools should be used and extended by end-users who are not experts in parallel discrete-event simulation, and whose primary goal is to describe and model their particular application domain, not rebuild a customized simulator for it.

SSF owes its existence to TeD, the pioneering modeling API prototype that made the critical leap from low-level, event-oriented simulation to high-level, process-oriented hierarchical modeling. The first implementations of TeD, atop Georgia Time Warp, and using CPP as the external language, have successfully shown that modelers can take advantage of parallel resources to run very large, complex simulations. TeD, in turn, has been designed to concisely express the key modeling concepts identified over the years of development of various hardware design languages, such as VHDL, SDL or Verilog.

The project stresses the necessity of verification of software design and implementation in realistically large and detailed modeling projects utilizing the modeling environment and simulators built by the S3 partners. This *co–design* process, involving networking engineers and graduate students, provides real life tests, exposes weaknesses, and feeds back corrections and new ideas to simulator and tool developers. The leading theme is versatility; thus we cover the extremes of the modeling problems: wireless networks, wired ATM internetworks, the global Internet, and the battlefield scenario–driven military communications models. In this way we gain confidence in the power of our approach to usable, versatile, scalable high performance simulations of telecommunications systems.