

1992 ANNUAL REPORT

International Computer Science Institute

February 1993

Contents

1	Research Results	1
1.1	Applications of Massive Parallelism	1
1.1.1	Knowledge Representation	1
1.1.2	Vision and Geometric Structures	1
1.1.3	Learning: Miniature Language Acquisition	1
1.1.4	Tools	1
1.2	Theory of Computation	2
1.2.1	Parallel Computation	2
1.2.2	Computational Learning Theory	3
1.2.3	Randomization as a Computational Resource	4
1.2.4	Foundations of Complexity Theory	6
1.2.5	Design and Analysis of Algorithms	8
1.3	Realization of Massively Parallel Systems	10
1.3.1	Systems Design	10
1.3.2	IC Design	11
1.3.3	Applications Work - Speech Recognition	12
1.3.4	Architectural Studies	14
1.3.5	Simulators	14
1.4	Very Large Distributed Systems and High-Speed Networking	14
1.4.1	DASH Group	14
1.4.2	Tenet Group	14
2	Project status compared to the original research plan	23
3	Prospects to reach original research goals	23
4	Impact of third party R&D results on the project	23
5	Changes and modifications in the research goals	23
6	Inventions, Patent applications, etc.	24
7	Bibliography	25

1 Research Results

1.1 Applications of Massive Parallelism

1.1.1 Knowledge Representation

A major achievement of 1992 was the completion of the joint project with Nigel Goddard at Carnegie Mellon [1] on connectionist models for the perception of moving light displays. A paper on the foundations of our approach to knowledge representation was published [2]. We were able to complete additional work in this area because of the addition of Gerhard Brewka, Marco Dorigo, Bernd Fritzke, Dan Jurafsky, and Matthias Kaiser as post-doctoral fellows. Current efforts in this area are focused on the representational and inference capabilities of structured connectionist models. Much of the work is connected to natural language semantics and the Miniature Language Acquisition task described below.

1.1.2 Vision and Geometric Structures

This work has proceeded according to plan with good progress being made on these very difficult problems. Steven Omohundro has completed the first basic paper on ‘Best-First Model Merging for Dynamic Learning and Recognition’ [3] and is applying this to concrete problems in the recognition of human gestures. This work is being aided by an Italian post-doc, Alessandro Verri and has resulted in a first experimental paper on recognizing hand shapes [4]. Much of the theoretical work on geometric structures has been codified in the form of Sather libraries and will be distributed in 1993. Work on general connectionist models for vision and space was mostly in connection with the MLA project described below.

1.1.3 Learning: Miniature Language Acquisition

We have been investigating the particular task (which we call L_0) of language acquisition in the domain of simple two-dimensional scenes. The MLA task requires that the system learn equivalent fragments of any natural language. We have been working on the L_0 task for about three years and have had a number of revealing experiences. The task has already changed the way we think about representation and learning in vision and language.

In 1992, the most important result was the completion of Terry Regier’s dissertation [5]. This system successfully learns simple static and motion spatial terms for a wide range of languages and is already receiving wide attention. There are several related projects that are continuing this work. Also, Andreas Stolcke completed his dissertation qualifying exam and published a paper on connectionist tree representation [6]. In the second half of the year Stolcke and Steve Omohundro published a significant new paper on learning stochastic grammars using minimum-description-length techniques [7].

1.1.4 Tools

We are developing several tools in support of the above research efforts. The components of these systems are very modular and it is natural to develop a “toolkit” which

would allow researchers to mix and match modules at a high level. As planned, the Sather programming system has been exported to many sites and a second version (1.0) is designed and now being implemented [8, 9]. A new parallel version has been implemented for the CM-5 computer and is being evaluated [10]. A major revision of our connectionist simulation tool, ICSIM, has been completed and documented [11]. Plans for the next year include converting the parallel Sather to the 1.0 version and constructing a parallel version of the ICSIM simulator. These efforts are also part of the CNS-1 software development discussed in Section 1.3.1.

1.2 Theory of Computation

1.2.1 Parallel Computation

During the next decade massively parallel computation will be dominated by machines consisting of a number of full-scale microprocessors, each with its own local memory, communicating by point-to-point messages through an interconnection network. Richard Karp has continued working with Professors David Patterson and David Culler of UC Berkeley and a group of students on the development of a model of parallel computation suitable for the design of efficient algorithms that can be ported to a wide class of distributed-memory machines [46]. An algorithm for the Fast Fourier Transform was analyzed within the model and programmed for the CM-5; the model's predictions proved generally accurate, although certain observed effects related to caching were not predicted. Within the model, optimal algorithms for summing N numbers and for simulating broadcast using point-to-point messages were developed. Among the broadcast problems considered, the most challenging problem, and one that is still only partially solved, is to minimize message delay in the case when a source generates a long stream of messages to be broadcast to all the processors in a P -processor system. Algorithms for sorting, connected components and LU decomposition have also been analyzed within the model.

Marek Karpinski and colleagues designed the first efficient deterministic NC algorithm for computing a Maximal Independent Set in a hypergraph of dimension 3. The algorithm runs in $O(\log^4 n)$ parallel time and $n + m$ processors on EREW PRAM and is optimal up to a polylogarithmic factor [47]. The algorithm improves on earlier algorithms of Dahlhaus-Karpinski, 1989, and Kelsen, 1990.

While visiting ICSI, Friedhelm Meyer auf der Heide worked with Richard Karp and Michael Luby to develop a randomized simulation of a $n \log \log(n) \log^*(n)$ -processor shared memory machine (PRAM) on an n -processor distributed memory machine (DMM) with optimal expected delay $O(\log \log(n) \log^*(n))$ per step of simulation [48]. The time bound for the delay is guaranteed with overwhelming probability. The algorithm is based on hashing and uses a novel simulation scheme. The best previous simulations use a simpler scheme based on hashing and have much larger expected delay: $\Theta(\log(n)/\log \log(n))$ for the simulation of an n -processor PRAM on an n -processor DMM, and $\Theta(\log(n))$ in the case where the simulation preserves the processor-time product.

In the context of parallel data structures and algorithms, Hans Rohnert and colleagues wrote sample applications for Travelling Salesman, Gröbner Basis Reduction and Finite Element Analysis. For the development of pSather they learned from these

experiments that the expressibility of pSather is well suited to describe coarse grain parallelism. Also, the speedup of the simple Travelling Salesman was very good.

Another goal for Rohnert was to investigate how hard it is to code PRAM algorithms in an asynchronous fashion on a distributed memory machine. Taking the example of parallel prefix, he and his colleagues came to the conclusion that this can be done in a natural way, though the actual coding takes quite some time.

While visiting ICSI, Noam Nisan worked with Michael Luby to develop a fast parallel approximation algorithm for the positive linear programming optimization problem, i.e., the special case of the linear programming optimization problem where the input constraint matrix and constraint vector consist entirely of positive entries [49]. They introduce an algorithm that takes as input the description of a problem and an error parameter ϵ and produces both a primal and a dual feasible solution such that the values of these two solutions are within a multiplicative factor of $1 + \epsilon$ of each other, and hence both these solutions have a relative error within ϵ of an optimal solution. Their algorithm can be implemented on a parallel machine using $O(N)$ processors with a running time polynomial in $\log(N)/\epsilon$, where N is the total number of non-zero entries in the constraint matrix. The algorithm is elementary and has a simple parallel implementation.

The outcome of a chess game is not predetermined but is a random event; the probability that the better of the two players wins depends on their relative strengths. Dick Karp, Claire Kenyon and two students have studied a model where the probability of error depends linearly on the difference of the values of the two players, and have constructed a randomized parallel tournament that ensures that with high probability the winner of the tournament is one of the best players, assuming a uniform distribution of values of the players entering the tournament. This is still work in progress.

1.2.2 Computational Learning Theory

In [50] Anna Morpurgo and her coworkers consider the problem of learnability in the distribution dependent framework. They introduce the notion of *polynomial uniform convergence* of relative frequencies to probabilities, give a necessary and sufficient condition for it, and study its relationship with polynomial-sample learnability. They show that, in contrast to what happens in the distribution independent setting, polynomial uniform convergence is a sufficient but not necessary condition for polynomial-sample learnability, and give an example from the boolean domain of a class of functions that is polynomial-sample learnable, but not polynomially uniform convergent.

Michael Kharitonov continued his investigations in the relationship between Cryptography and Computational Learning Theory. The paper “Cryptographic Lower Bounds for Learnability of Boolean Functions on the Uniform Distribution” was presented at COLT 92. The journal version of the paper “When Won’t Membership Queries Help?” (with Dana Angluin) was submitted for publication. Current research includes cryptographic hardness of Bayesian learning and separations of learning models.

The research of Marek Karpinski and his colleagues was conducted on computational complexity of learning boolean functions over a different basis. An efficient algorithm has been presented for exactly learning the class of thresholds read-once formulas. Using a generic transformation method of Angluin-Hellerstein-Karpinski, 1989, they were able to convert this algorithm into several other learning models [51].

Thorsten Werther continued his investigations in the computational complexity of learning sparse polynomials. He proves in [52] that the Vapnik-Chervonenkis (VC) dimension of t -sparse linear combinations of Chebyshev polynomials is exactly $2t$ in the halfplane $x \geq 1$, and is infinite in the strip $|x| \leq 1$. The first result implies distribution-free uniform learnability of this function class when the examples are drawn from the halfplane $x \geq 1$ and extends the results of Karpinski and Werther [53], [54] on learnability of sparse linear combinations over the standard base $\{x^i\}$. The proof involves a generalization of Descartes' Rule of Signs.

1.2.3 Randomization as a Computational Resource

For the randomized analysis of many algorithms only pairwise independence between the random variables is assumed. For the analysis of other randomized algorithms stronger randomness properties are needed.

Guy Even, Oded Goldreich, Michael Luby, Noam Nisan, and Boban Veličković describe constructions of sample spaces for n many-valued not necessarily uniformly distributed random variables [55]. Here, the probability distribution on n general m -valued random variables is described by a n by m *stochastic matrix* $\mathcal{P} = \{p_{i,v} : i \in \{1, \dots, n\}, v \in \{1, \dots, m\}\}$. The (i, v) -entry $p_{i,v}$ specifies the probability that the i^{th} random variable should take on value v . For all values of $l \in \{1, \dots, n\}$, for all $I = \langle i_1, \dots, i_l \rangle$, where $1 \leq i_1 < \dots < i_l \leq n$, and for all $V = \langle v_1, \dots, v_l \rangle \in \{1, \dots, m\}^l$, let $p_{I,V} = \prod_{j=1}^l p_{i_j, v_j}$ be the probability that the subsequence of random variables indexed by I should take on value V if the random variables were totally independent. We say sample space S is a (k, ϵ) -*approximation* for \mathcal{P} if for any subsequence I of size $l \leq k$ and for any set of possible values $V \in \{1, \dots, m\}^l$, $|\Pr_S[x_I = V] - p_{I,V}| \leq \epsilon$. S is an ϵ -*approximation* if this statement is true with no restriction on the size of I and S is a k -*wise independent approximation* if this statement is true with $\epsilon = 0$.

The first construction yields a sample space that is a (k, ϵ) -approximation, where the size of the sample space is polynomial in $\log(n)$, 2^k and $1/\epsilon$. The second and third constructions for n general random variables yield sample spaces that are ϵ -approximations for general \mathcal{P} , where the size of the sample space is polynomial in $(n/\epsilon)^{\log(1/\epsilon)}$ for the second construction and polynomial in $(n/\epsilon)^{\log(n)}$ for the third. In contrast, the previous bound on the sample space size, implicit in the classical work on discrepancy theory is polynomial in n^n/ϵ . For interesting cases of n and ϵ , i.e., when $1/\epsilon$ is polynomial in n , the results presented are a dramatic improvement over what was previously known.

While visiting ICSI, Nati Linial, Michael Saks and David Zuckerman have worked with Michael Luby to develop a polynomial time deterministic algorithm for the problem of finding a hitting set for combinatorial rectangles in high dimensional space [56]. This problem has been studied extensively in the classical case when the dimension is a fixed constant. One natural reason for considering the problem when the dimension is an input parameter is that a solution yields an explicit polynomial size sample space for a large number of random variables (the number of variables is the dimension of the space) that approximates the independent distribution on the random variables.

The input parameters to the problem are d , ϵ and m , where d is the dimension, m is the number of values possible in each dimension and ϵ is a minimum volume parameter.

A combinatorial rectangle R is defined by $R_1 \subset \{1, \dots, m\}, \dots, R_d \subset \{1, \dots, m\}$, where points in R consist of all d -tuples r_1, \dots, r_d such that $r_i \in R_i$ for all $i = 1, \dots, d$. The volume of R , $\text{vol}(R)$, is simply $\prod_i (|R_i|/m)$. We can deterministically and efficiently construct a set $S = \{s = \langle s_1, \dots, s_d \rangle : s_i \in \{1, \dots, m\}\}$, where $|S|$ is polynomial in $\log(d)$, m and $1/\epsilon$, with the following property. For all combinatorial rectangles R , $\text{vol}(R) \geq \epsilon$ implies that $S \cap R \neq \emptyset$.

Michael Luby, Boban Veličković and Avi Wigderson developed deterministic algorithms which, for a given depth-2 circuit F , approximate the probability that on a random input F outputs a specific value α [57]. Their approach gives an algorithm which for a given $GF[2]$ multivariate polynomial p and given $\epsilon > 0$ approximates the number of zeros of p within a multiplicative factor $1 + \epsilon$. The algorithm is deterministic and runs in time $\exp(\exp(O(\sqrt{\log(n/\epsilon)})))$, where n is the size of the circuit. They also obtain an algorithm which, given a DNF formula F and $\epsilon > 0$, approximates the number of satisfying assignments of F within a factor of $1 + \epsilon$ and runs in time $\exp(O((\log(n/\epsilon))^4))$.

Recently, Karger introduced a fast randomized parallel algorithm to find a minimum cut in an n -vertex graph. His algorithm uses polynomial in n random bits. In recent work in progress, Michael Luby, Moni Naor and Seffi Naor developed a randomized version of this algorithm that uses $O(\log^2(n))$ random bits, thus partially derandomizing the original algorithm [58]. They break up the original algorithm of Karger into $O(\log(n))$ phases, and then use a probabilistic analysis for each phase that assumes only pairwise independence between the random variables. The result follows because there are known explicit constructions of sample spaces in which the random variables are pairwise independent such that it takes $O(\log(n))$ random bits to choose a random point in the sample space.

In a development that is an outgrowth of the monograph in progress by Michael Luby titled “Pseudorandomness and Applications” [59], Amir Herzberg and Michael Luby have formalized the notion of public randomness in cryptography, and demonstrate its importance to the analysis and design of cryptographic primitives [60]. Many uses are known for public random bits in cryptography. Previous works considered the public random bits as a part of the input, and security was parameterized in terms of the total length of the input. Here security is parameterized solely in terms of the length of the private input, and treats the public random bits as a separate resource. This separation allows one to independently address the important issues of how much security is preserved by a reduction and how many public random bits are used in the reduction.

In this work, Herzberg and Luby study reductions from weak one-way permutations to one-way permutations. Based on the new definitions, they present simple reductions with security preserving properties stronger than previously known. The simplest of the reductions uses a large amount of public randomness. They present a series of increasingly intricate reductions with the goal of minimizing the number of public random bits used in the reduction. The final reduction is a slight modification of a reduction described by Goldreich et al. (Goldreich, Hastad, Impagliazzo, Levin and Venkatasawan). This reduction has stronger preserving properties than the reduction described in Goldreich et al. (even under the new definitions) and, as is the case for their reduction, uses only a linear number of public random bits.

In continuing work with Yuri Rabinovich (who visited ICSI in October) and Avi Wigderson of the Hebrew University in Jerusalem, Alistair Sinclair has been investigating a class of random processes known as Quadratic Dynamical Systems, which have many computational applications. These processes model any system which evolves via random pairwise interactions that are locally reversible: examples include Boltzmann's ideal gas model, simple models in population genetics, and the mating operation in Genetic Algorithms for combinatorial optimization. In order to design and analyze algorithms based on such systems, it is necessary to understand their equilibrium behavior and the rate of convergence to equilibrium. In [61], a theoretical framework for the analysis of these systems is presented, and the first convergence rate results established for a non-trivial system defined over combinatorial structures. The authors are currently extending these ideas with the aim of providing general tools for the quantitative analysis of Boltzmann-type systems and Genetic Algorithms.

The computational problem of determining the number of non-overlapping arrangements of dominoes on a d -dimensional lattice has occupied mathematicians and theoretical physicists for several decades. Apart from its inherent combinatorial interest, it has applications to the theory of monomer-dimer systems and the cell-cluster theory of the liquid state. It seems unlikely that the problem can be solved *exactly* in polynomial time in dimensions higher than two, but in recent work Claire Kenyon (visiting ICSI in Fall), Dana Randall of UC Berkeley, and Alistair Sinclair have shown the existence of the first polynomial-time randomized algorithm that computes the number in arbitrary dimension *approximately*, to within any specified accuracy with very high probability [62]. The algorithm draws on earlier work by Sinclair and Mark Jerrum on rapidly mixing Markov chains.

Marek Karpinski and colleagues have been studying the impact of randomization on algebraic decision complexity. Among other problems they study the determinantal test complexity. Examples have been given where the randomization dramatically reduces the decision complexity. A lower bound has also been proven, showing that in some cases randomization cannot help (much) [63].

They prove that depth d threshold circuits of size $n^{O(1)}$ can be simulated by uniform majority circuits of depth $d + 1$ and size $n^{O(\log n)}$. The proof method involves the derandomization of the construction of Goldmann, Hastad, and Razborov (1992), and the construction of a pseudorandom generator that can fool weighted sums of boolean variables (for the small weights) [64].

1.2.4 Foundations of Complexity Theory

Continued progress has been made on the book/research monograph "Theory of Computation and Complexity over the Reals" by Lenore Blum and colleagues. This work integrates new results of Lenore Blum, Michael Shub, and Steve Smale, as well as other researchers (including James Renegar, Salil P. Vadhan, Chi Tat Chong, Christian Michaux) in continuous complexity theory with mainstream work in numerical analysis and dynamical systems, classical logic and complexity theory. A focus of research in this area has been to understand the inherent relationship between the condition of a problem and its complexity. ICSI external fellow James Renegar has investigated this relationship for the linear programming problem. ICSI visitor Michael Shub, together

with Steve Smale, has investigated this relationship for the problem of solving systems of polynomials, in particular in the context of Bezout's Theorem. Crucial to this development is the formulation of natural measures of condition. Related work, primarily of graduate student Doug Priest, incorporates round-off error into the BSS model of computation.

L. Blum, and S. Smale have begun to investigate entropic measures of complexity in the BSS model of computation. The "topological" entropy is a measure of the proliferation of computation paths in a machine and relates to its branching complexity. It is lower bounded by the more classical Shannon type "information-theoretic" entropy. Blum and Smale are investigating how these notions of entropy relate to the computational complexity of such machines.

Karl Aberer and Bruno Codenotti continued their work on models suited to study complexity-error tradeoffs for numerical problems. The tools are introduced in [65] and [66] and include circuit models for approximate and adaptive computations, models for error and adversaries, as well as algebraic models (see also [67] and [68]).

It was shown that for oblivious algorithms there exist lower bounds on the error introduced by an approximate computation. Adaptive algorithms can avoid this error. An example studied in particular was the floating point addition of n numbers, where a new adaptive algorithm could be given, which avoids errors completely, computing the exact sum in $O(n \log n)$ steps.

In the framework of these models also the relative power of discontinuous operations, namely division and conditional operations, in approximate computations was studied [69]. More precisely, the simulation of conditional operations by rational functions showed that adaptive algorithms can be simulated efficiently by oblivious algorithms, except for the ill-conditioned inputs. It was also shown that the simulations cannot be performed efficiently by using only polynomial operations. Together with Luciano Margara and Giovanni Resta of IEI-CNR, Pisa, Aberer and Codenotti also started a book on complexity theory, to be titled "The Many Facets of Complexity." This book will also include, among other topics, the recent work at ICSI of Karl Aberer and Bruno Codenotti, as well as other ICSI-affiliated researchers (Lenore Blum, Michael Shub, Steve Smale, Klaus Weihrauch).

The NP-completeness of a decision problem rules out the possibility of finding an optimal solution of the corresponding optimization problem in polynomial time unless $P = NP$. It does not exclude, however, the possibility of finding efficient algorithms that produce approximate solutions. In fact, for many optimization problems with NP-complete decision problems, there are simple and efficient algorithms that produce solutions differing from optimal solutions by at most a constant factor. For some problems, there even exist so-called polynomial-time approximation schemes, which produce approximate solutions to any desired degree of accuracy. For other problems, notably the TSP, there do not exist efficient approximations unless $P = NP$. The "structural" reasons for the different approximation properties of NP optimization problems are not yet sufficiently understood.

Erich Grädel, together with Thomas Behrendt (University of Basel) and Kevin Compton (University of Michigan, Ann Arbor and INRIA, Rocquencourt), investigated some of these "structural" reasons. Their work extends earlier research of Papadimitriou and Yannakakis, as well as that of Kolaitis and Thakur.

1.2.5 Design and Analysis of Algorithms

Richard Karp and Farid Alizadeh have been working with UC Berkeley students Lee Newberg and Deborah Weissner and with Geoffrey Zweig to develop and implement algorithms for the physical mapping problem. They consider the case where, for each clone, and for each of a number of short DNA sequences called *probes*, the problem data indicates whether the probe hybridizes to (i.e., occurs on) the clone. Given this data, they seek to determine the most likely pattern of overlap of the clones.

For the case where the probes occur according to a Poisson process they have developed and implemented an algorithm that has performed successfully on simulated data [70]. The algorithm employs certain hillclimbing heuristics similar to those often used for the travelling-salesman problem. They are now investigating whether a more powerful heuristic called Divide-and-Merge, invented by Zweig and already applied successfully to the travelling-salesman problem, will improve the performance of the algorithm.

Recently, physical mapping efforts within the human genome project have focused on unique probes; i.e., probes that occur only once along the extent of the DNA. The ICSI group has developed a heuristic for constructing physical maps from unique probe data and applied it both to artificial data and to data from human chromosome 21. The heuristic has given excellent results, and has proven to be quite robust, working well even when the data contains many errors. We are also studying a version of the physical mapping problem in which the extraction of unique probes from the DNA and measurement of their hybridization to clones is carried out in an on-line fashion, so that the choice of the location of each probe is determined by the outcomes of the earlier hybridization experiments.

A common experimental technique is to cut a DNA molecule with a restriction enzyme that cuts selectively at certain *restriction sites*, and to then reconstruct the locations of the restriction sites from noisy measurements of the distances between all pairs of restriction sites. For one version of this problem, known as the Probed Partial Digest Problem, Lee Newberg and Richard Karp have developed and implemented an efficient algorithm.

Richard Karp has investigated the following generalization of binary search: minimize the number of function evaluations to determine a nondecreasing integer-valued function throughout its domain. His paper [71] determines the worst-case and expected complexity of sequential and parallel deterministic algorithms for this problem, and also investigates the complexity of randomized algorithms for the problem.

Continuing their investigation of on-line investment problems [72] Richard Karp and Ran El-Yaniv of the University of Toronto have considered a class of *on-line replacement* problems. For such a problem, the basic question is when to switch from one activity, investment or facility to another more rewarding one, when there is a *changeover cost* associated with making the switch. One example is the problem of when to refinance a mortgage.

The quality of an on-line policy for such a problem is measured by its *competitive ratio*: the worst-case ratio between its cost and that of a clairvoyant policy. For a certain abstraction of the mortgage refinancing problem El-Yaniv and Karp have determined a policy with optimal competitive ratio.

The paper [73] is a survey of the growing body of research on the competitive analysis

of on-line algorithms.

In this framework Alberto Marchetti-Spaccamela, Umberto Nanni, and a colleague have considered the *Weighted List Update Problem* [74], where the cost of accessing an item depends on the item itself. Two heuristics are presented for this problem, and both are shown to be 2-competitive against a *lazy adversary*.

Some problems in incremental compilation can be naturally reduced to dynamic graph algorithms. Specifically these are reachability from a designated start node, cyclicity test, topological order, parentship and common descendant (the latter three ones are of particular relevance for compiling object-oriented programs). Several algorithms were devised and are being written down in the process of checking which of these methods can be employed during the construction of the new Sather 1.0 Compiler. In particular Alberto Marchetti-Spaccamela, Umberto Nanni, and Hans Rohnert overview the current best solutions for several of these problems, and propose a new algorithm for the dynamic maintenance of a topological order [75].

Marek Karpinski and colleagues analyzed the computational complexity of interpolating real algebraic functions given by a black box for their evaluations, and gave the first algorithm for interpolating sparse algebraic functions from a black box [76]. They display an existence of short proofs for the nondivisibility of two sparse multivariate polynomials under the Extended Riemann Hypothesis (ERH). They also prove that the same divisibility problem for rational functions given by a black box belongs to the parallel class NC^2 , provided the degree of some sparse representation of it is known [77].

Michael Kharitonov completed the work for the First DIMACS Implementation Challenge. The paper “On Implementing Scaling Push-Relabel Algorithms for the Minimum-Cost Flow Problem” (with A. Goldberg) was submitted for publication by DIMACS.

The *All Pairs Shortest Paths Problem* in a directed graph $G = (V, E)$ where edges may be dynamically inserted or have their cost decreased is considered in [78] by Giorgio Ausiello, Alberto Marchetti-Spaccamela, Umberto Nanni, and a colleague. The total time required to maintain the data structure under any sequence of edge insertions and edge cost decreases is $O(Cn^3 \log(nC))$ in the worst case, where n is the total number of vertices in G and the edge weights range in $[1 \dots C]$. The same bounds can be achieved for the problem of maintaining on-line longest paths in directed acyclic graphs. These algorithms improve previously known algorithms and are only a logarithmic factor away from the best possible bounds.

Directed *hypergraphs* are a generalization of directed graphs, introduced to represent the functional relation inherent in the structure of several computational problems. Giorgio Ausiello, Umberto Nanni and a colleague propose in [79] a quite general framework to deal with optimization problems on hypergraphs, either determining their intractability, or providing algorithms to find and maintain optimal *hyperpaths*.

Directed hypergraphs have been proposed to model Conflict-Free Petri Nets. Umberto Nanni and colleagues considered in [80] several problems for this class of nets (place reachability, liveness, and boundedness) proposing linear time algorithms which dramatically improve the best previous known results for these problems. In the incremental solution of these problems queries are answered in constant time, and the total cost for all the updates is still linear in the size of the final net.

Rotations are an operation for transforming binary trees which is used as a primitive for defining splay trees. In 1986, Sleator, Tarjan and Thurston proved that the maximum distance between two trees using rotations is exactly $2n - 6$ if n , the number of nodes of the trees, is large enough. Their surprising proof was based on hyperbolic volume arguments. Claire Kenyon and Bill Thurston have worked on an elementary proof of that result, using max-flow-min-cut arguments. This work is still in progress.

1.3 Realization of Massively Parallel Systems

1.3.1 Systems Design

In 1992 we designed a memory upgrade for the RAP [81] which quadrupled the per-node SRAM, and worked on simplifying and completing the user documentation for the RAP, as well as consolidating the software libraries. This system was incorporated in a new manufacturing run of RAP machines, primarily aimed at a number of European laboratories (Cambridge in the U.K., INESC in Portugal, and Lernout & Hauspie in Belgium), but also at Siemens in Princeton (Siemens - Munich already has RAP boards). The reorganization of the software is an important step in the further development of the Synthetic Perceptron Trainer (SPERT) architecture, which is being incorporated into a system for training our connectionist speech recognition systems [82].

We designed a shared interface called the Common Server Interface (or CSI). Object classes (such as matrix and vector) that insulate the user from the hardware configuration are built on top of the CSI class. The CSI class defines the interface that is inherited by classes such as CSLSPARC, CSLRAP and CSLSPERT. These three classes contain the implementation of CSI for each of the three hardware configurations. Multiple CSLRAP or CSLSPERT objects can be created allowing the host to command more than one RAP or SPERT system.

We completed the paper design of SPERT, as well as the physical fabrication and testing of one of the SPERT data paths (see IC work below) [82][83][84][85].

Much of our effort as a group went into directed studies concerning the RAP's successor, the Connectionist Network Simulator (CNS-1). A preliminary design had been proposed in 1991, but in 1992 the project began in earnest, beginning with an active questioning of nearly all conclusions from the 1991 proposal document. This work was done in close cooperation with the Applications group. The discussions culminated in an open design review, attended by many experts in the areas of concern (e.g. David Patterson, Dan Hammerstrom, Dick Lyon). As a result of both our own discussions and reflection on the design review, we made a number of decisions that differed somewhat from the original conception of the CNS-1. Still in the design is the basic node idea (a scalar RISC core tightly coupled with a group of arithmetic datapaths optimized for multiply-accumulate operations). Some of the important decisions made this year were:

1. We are planning on the use of Rambus-based DRAM or RDRAM. This is a new form of dynamic memory that incorporates low-voltage signaling levels and a high-speed packet-based processor interface, capable of a peak transfer rate of 500 MB/sec.
2. Topology: we moved from a ring-of-rings to an end-connected mesh (2 edges connected, as opposed to the 4 of a torus), referred to as a *cylinder* or a *barrel*. This

change provided us with better near-neighbor communication for computation such as local receptive fields for early vision, while continuing to give a good structure for other architectures such as the feedforward networks that we have incorporated in our speech research.

3. Operating voltage: to reduce power requirements (and corresponding cooling) we decided to run the system at 3.3 Volts instead of 5V. An investigation of the process parameters for the candidate fabrication processes showed that this was feasible.
4. Scalar processor: we have decided to use an industry standard instruction set architecture (ISA) for the scalar processor within the Torrent chip (the computational node chip for the CNS-1). Although this decision directly only affects about 15% of the chip area (making it slightly more complex to design), the positive repercussions are extensive for software development. We have chosen the MIPS R3000 ISA as our standard. By using this standard, we will be able to take advantage of suitable commercially available and public domain system software, tools, and application libraries. In addition, machines executing the R3000 instruction set are widely available and will be used as development platforms throughout the life of the CNS-1.

Because of this final decision, we restructured the SPERT design to also incorporate the MIPS R3000 ISA. Since this implied rather fundamental architectural changes (e.g., since this ISA used a 32-bit instruction word, our design could no longer be VLIW), we decided to rename the new chip rather than continue to call it SPERT. This chip will be called T0 (T-zero), as it can be viewed as the initial implementation of the Torrent processor. Like that processor, it will incorporate a vector architecture. Unlike Torrent, it will fit into a SPERT-card S-bus system, running at a lower speed (50 MHz), using a single bank of external SRAM and requiring no network interface.

Work on the software and applications (for both the SPERT board and for CNS-1) continued in collaboration with the Applications Group.

1.3.2 IC Design

We have finished a significant amount of custom IC design. The major developments are:

1. Design path established - we have converged upon a design and simulation path for macrocells, using CAzM¹ (a table-driven approximation to SPICE), Viewlogic² (a schematic editor), IRSIM (a switch-level simulator), Magic (a layout editor), and Gemini, a netlist comparator used for layout vs. schematic (LVS) comparisons.
2. Implementation strategies formulated - we have selected a clocking methodology and floorplan for our macrocells that allow them to be easily used together and ensure a high level of performance. We are using True Single-Phase Clocking

¹CAzM is a product of the Microelectronics Center of North Carolina (MCNC)

²Viewlogic is a registered trademark of Viewlogic Systems, Inc.

(TSPC) [86] which has several advantages over the more common two-phase non-overlapping clocking scheme. In particular, TSPC is more area-efficient and is somewhat faster, since no “dead time” is required to separate clock phases. Additionally, this scheme reduces power requirements and skew. Finally, we have chosen layer assignments and consistent floorplan conventions for macrocells to permit tiling components together.

3. Datapath completed - we have completed the design, fabrication, and testing of a datapath test chip that we call SQUIRT. This datapath incorporates the arithmetic modules that we have built up for connectionist algorithms. The chip confirmed the correct operation of a number of modules, and also pinpointed errors that have since been corrected.
4. Instruction cache completed - we completed the design, fabrication, and testing of a direct-mapped instruction cache test chip. This cache will be used for the completion of the SPERT project and also for CNS-1.
5. Pads and test structures - we have fabricated and tested a number of pad types and confirmed their correct operation and have gathered data about their operating characteristics. We also have determined some electrical characteristics for neighboring wires that were not available from MOSIS.
6. ISA design for the Torrent and T0 designs - as noted above, in order to accommodate the new decision to conform to an R3000 scalar processor, we had to convert from the original VLIW SIMD datapaths to a vector structure. This necessitated a number of design changes. These have initially been done at the level of an ISA specification.

1.3.3 Applications Work - Speech Recognition

The RASTA robust speech analysis approach was tested further and extended during this period [87]. The basic principle of this technique was to do some band-pass filtering of the power spectral trajectory (over time) in some nonlinear domain (e.g. logarithm). In its original form (using the log domain), this technique provided robust recognition despite multiplicative changes in the overall spectrum (e.g. from a change in microphone or telephone channel). In one experiment in 1992, we took continuous sentences recorded simultaneously over eight different microphones and reduced the recognition error (using a system trained on one of the microphones under controlled conditions) by roughly one-half [88]. Additionally, we began a project to extend the technique to robustness to the combined effect of multiplicative spectral change and additive noise. Preliminary results indicate that this could be accomplished without a long noise spectrum estimation phase [89][90]. This was done by developing a new form of RASTA in which the filtering domain was a more complex function, one which was roughly logarithmic for large spectral values and low noise, and roughly linear for small spectral values and/or high noise. This was combined with a multiple-recognizer approach in which the correct nonlinearity was applied based on a gross estimate of overall noise level. For both an isolated speech and a continuous speech recognition example, this approach appeared

capable of providing at least as good a performance as one could get by using training data from the degraded testing condition.

In a separate project, an approach was developed for the connectionist mapping of LPC parameters from noisy speech to the corresponding parameters for clean speech. This was tested on an isolated word recognition task and appeared to be moderately robust over some changes in noise sample and noise amplitude [91] [92].

We have continued our research in connectionist phonetic probability estimation [93]. Recently this work has been extended in collaboration with researchers at SRI International and Stanford [94] [95]. Context-dependent probability estimators, discussed in [96], are smoothed with context-independent estimators using cross-validation [97]. Architectures implementing this approach have yielded significant improvements in recognition performance. We also showed that the probabilistic interpretation of networks trained for classification could be used as a general technique for factoring large networks into smaller ones without strong assumptions such as statistical independence [98]. Additionally, new training has been implemented using random pattern presentation and pre-setting of the unit biases to negative values that approximate the prior probabilities of each corresponding class. These modifications, along with simply increasing the size of the nets trained, have led to significant improvements in the context-independent networks. Finally, the probabilities obtained from these networks have been smoothed together with probabilities obtained from SRI's tied-Gaussian mixture estimators, and the resulting recognizer obtained the best score for February 1991 standard Resource Management evaluation test set, as measured by the National Institute of Standards and Technology (NIST) [99] for the DARPA neural network speech program in early 1992.

New gender-dependent enhancements to our recognizer have improved performance further, and now require fewer parameters than the earlier system [100]. Recent experiments have also shown that similar results can be obtained for speaker cluster classification in which the clusters are obtained by unsupervised learning [101].

Work has proceeded on the voice-interactive Berkeley Restaurant guide. In-house recordings were made using a Wizard or PNAMBC (Pay No Attention to the Man Behind the Curtain) system, or one in which the recognizer is replaced by a human listener without the knowledge of the system user. This subterfuge is used to attempt to influence the system user to structure his speech as he would to the final machine system. Sufficient speech was collected to bootstrap the first automatic system that we are currently completing. Additionally, we have worked on developing self-organized speech sub-word units to assist us in developing automatically-generated phonological models for the restaurant vocabulary. Some new approaches to this problem, similar in spirit to the IBM "fenones", are being tried [102] [103].

Additional work on the restaurant project has been the development of methods for the detection and modeling of foreign accents (in particular, German vs. native American) in English. We are using a combination of frame-level connectionist approaches and some study of syntax.

In collaboration with the Applications Group, we are also working on incorporating a probabilistic context-free grammar language model for the Restaurant project, including significant use of semantic knowledge in the grammar. We are working on approaches to tight coupling between natural language modules and the acoustic recognizer. Also

in collaboration with the Applications Group, we have incorporated the HMM learning algorithm of Stolcke and Omohundro [7] to provide a bottom-up learning of a multiple-pronunciation lexicon.

1.3.4 Architectural Studies

The work in this section is now merged with the Systems Design section.

1.3.5 Simulators

The ICSIM simulator work has explicitly migrated to the Applications Group, as has the GUI work.

CLONES has been put into daily use by a growing number of researchers, and so has been modified in many small ways to accommodate the users' needs. In addition, we are nearing completion of a new version of CLONES that provides a simpler approach to plugging together different network elements, called Boxes-Of-Boxes (BOB). We expect this to be complete by the end of the first quarter of 1993.

1.4 Very Large Distributed Systems and High-Speed Networking

1.4.1 DASH Group

The project leader of the DASH project, David P. Anderson, left UCB and ICSI. This required merging the DASH project with the Tenet Group. As can be seen in the ICSI Annual Report 1991, most major DASH project portions were completed in 1991. The remaining work concerning multimedia support in distributed systems will now be conducted under the auspices of the Tenet Group. This also makes sense from a technical/scientific point of view since multimedia applications are of great interest to the Tenet Group. The work in this area is described in the "Multimedia Support" section below.

1.4.2 Tenet Group

The Tenet Group is continuing its activity toward the implementation of the Tenet real-time protocol suite for guaranteed communication performance on high-speed networks. Prototype versions of the Real-Time Internet Protocol (RTIP), the Real-time Message Transport Protocol (RMTP), and the Real-time Channel Administration Protocol (RCAP) have been developed and are currently under testing on a local high-speed network testbed. The design and coding of the Continuous-Media Transport Protocol (CMTP) has been completed; its embedding into a Unix kernel is currently under testing.

Other activities of the Tenet Group during 1992 have been concerned with support for multimedia, performance analysis of multimedia systems, analysis of Internet packet traffic, real-time and distributed operating systems requirements, and models for client-network quality-of-service negotiation.

A set of experiments to be performed in 1993 has been designed for two major multi-institution research projects the Tenet Group is participating in: BLANCA and Sequoia 2000. The BLANCA project is part of the Gigabit Testbed Initiative, which is

sponsored by the Corporation for National Research Initiatives and supported by the National Science Foundation and by the Defense Advanced Research Projects Agency. The BLANCA project officially started in May 1990 and currently includes AT&T Bell Laboratories, the universities of California, Illinois, and Wisconsin, Pacific Bell, and a few research institutes, among which is ICSI. The goal of BLANCA is to design and build two wide-area ATM-based testbeds called Xunet 2 (at 45 Mbps) and Xunet 3 (at 622 Mbps), and demonstrate on them several applications requiring high-speed networking.

Sequoia 2000 aims at building high-speed storage and communication facilities for global-change scientists at remote locations to visualize and browse sequences of digitally represented satellite maps [104]. Project Sequoia 2000, the sequel to Project Athena, was awarded in July 1991 by the Digital Equipment Corporation to the University of California system; its network design and research activities are taking place at UC Berkeley and UC San Diego. The Sequoia network connects at T1 speed the Berkeley campus to the Santa Barbara, Los Angeles, and San Diego campuses, as well as to the Scripps Institution of Oceanography in San Diego and the State's Department of Water Resources near Sacramento. The network's bandwidth will be increased to 45 Mbps in June 1993.

Performance Control for Real-Time Services: Can end-to-end communication performance be guaranteed by a packet-switching internetwork? Domenico Ferrari has investigated this question by examining the feasibility of extending to an internetwork the Tenet approach to real-time communication service design. The conditions an internetwork must satisfy for the approach to be extendible to it have been investigated as well. These include conditions for the scheduling discipline to be used in the internetwork's nodes. The original Tenet approach applies to a network consisting of hosts, homogeneous nodes (or switches), and physical links connecting nodes and hosts in an arbitrary topology. The nodes are store-and-forward, and are scheduled by a multi-class version of the Earliest Due Date deadline-based policy. Ferrari showed that the extendibility conditions are quite broad, and this suggests that the Tenet approach may be used to establish and run real-time channels in a vast class of internetworks [105].

The problem of delay jitter control in high-speed networks has also been studied. Delay jitter is the variation of the delays with which packets travelling on a network connection reach their destination. For good quality of reception, continuous-media (video, audio, image) streams require that jitter be kept below a sufficiently small upper bound. Domenico Ferrari has proposed a distributed mechanism for controlling delay jitter in a packet-switching network [106]. The mechanism can be applied to an internetwork that satisfies certain conditions and can coexist with other schemes (or with the absence of schemes) for jitter control within the same network, the same node, and even the same real-time channel. The mechanism makes the distribution of buffer space requirements more uniform over a channel's route, and reduces by a non-negligible amount the total buffer space needed by a channel.

In the area of queuing techniques for real-time scheduling, Hui Zhang and Domenico Ferrari have proposed a new service discipline, called the Rate-Controlled Static-Priority (RCSP) queueing discipline, that can provide throughput, delay, delay jitter, and loss-free guarantees in a connection-oriented packet-switching network [107]. Previously proposed solutions were based on either a time-framing strategy or a sorted priority

queue mechanism. Time-framing schemes suffer from the dependencies that they introduce between the queueing delay and the granularity of bandwidth allocation; sorted priority queues may be difficult to implement. The RCSP queueing discipline decouples the allocation of delay and bandwidth by separating rate-control and delay-control functions in the design of the server.

The protocols composing the Tenet real-time protocol suite have been the object of an intense activity. The Real-time Channel Administration Protocol (RCAP) provides control and administration services for the Tenet real-time protocol suite, a suite of network and transport layer protocols for real-time communication. RCAP performs per-channel reservation of network resources based on worst-case analysis to provide guarantees on delay, jitter, and packet loss bounds. It uses a hierarchical approach to obtain these guarantees across a heterogeneous internetwork environment. Anindo Banerjea and Bruce Mah have recently completed a prototype of RCAP. Although the RCAP software has been completed and tested, it must be configured for proper operation in the network where it is to be used. In particular, the instance of the RCAP software running on each of the various network nodes needs some characterization of that node's performance. Examples of such performance parameters include the link speeds of each of the attached network interfaces and the packet processing delays incurred in the node itself. Towards this end, Banerjea and Mah are now working with Hui Zhang and Tom Fisher to determine the values of the necessary parameters, so that RCAP can be run on the Tenet Group's experimental FDDI testbed. The information and techniques gained during this effort will help determining these parameters for the different networks (for example, Xunet 3) where RCAP will eventually be installed.

Hui Zhang has implemented the Real-Time Internet Protocol (RTIP) and the Real-time Message Transport Protocol (RMTP) on HP 9000/700 and DEC 5000 workstations. RTIP is a network layer protocol that exists in both communicating hosts and network gateways. It provides an end-to-end, unreliable, in-order, guaranteed-performance packet service. The protocol implements scheduling mechanisms that allow packet forwarding and delivery in accordance to the performances specified by the users at connection time. RMTP is a transport layer protocol. It provides an end-to-end, unreliable, in-order, guaranteed-performance message service. RTIP and RMTP are also going to be implemented on SGI Iris-4D and Sun 4/280 workstations. The RMTP/RTIP protocol stack coexists with the TCP/IP protocol stack in the kernel, and provides the familiar socket interface to communication clients. Hui Zhang and Tom Fisher are measuring the performance of the RTIP and RMTP implementation. The measurements are performed in a local testbed consisting of a couple of DECstation 5000's connected by an FDDI ring. Preliminary results show that, with no extra loads, the throughputs and delays obtained by using RMTP/RTIP are comparable to those obtained by using raw IP or UDP/IP without checksums. When extra load (more IP traffic) is introduced, however, the performance of raw IP and UDP/IP declines considerably, while the performance of RMTP/RTIP is almost unaffected [108]. Another set of measurements are intended to identify the bottlenecks in the implementation and the architectural defects with respect to high speed networking in the current DECstation/Ultrix environment.

During 1992, facilities were set up for the segment of the Xunet 3 testbed between the University of California at Berkeley campus and the Lawrence Berkeley Laboratory. The fiber-optic links between the two sites have been installed and tested; UCB and LBL

researchers have begun experimenting with some of the fibers to carry FDDI network traffic. Various pieces of HIPPI network equipment have been procured and installed. A HIPPI Serial Extender from Broadband Communications Products has been tested over short distances in the laboratory. Two host interfaces for VME-based machines (in this case, Sun 4 workstations and SPARCstations) have been installed and tested in loopback configurations. Bruce Mah and Tom Fisher have begun porting the Tenet real-time protocol suite to the HIPPI network. This work consists of producing a working SunOS kernel from source code and then adapting and integrating into it the existing real-time protocol code. Of particular interest are the changes in the operating system and machine architecture (the prototype suite was developed on DECstation 5000's running DEC Ultrix, while the Xunet 3 hosts are SPARCstations and Sun 4's running SunOS). It is hoped that these differences, with respect to the protocol suite, will be minor (the networking code in the two operating systems is almost the same).

Integrated services networks are expected to provide a wide variety of services, which could consume widely differing amount of resources. Colin Parris, Srinivasan Keshav, and Domenico Ferrari have proposed a framework for pricing services in integrated networks, and studied the effects of pricing on user behavior and network performance [109]. They have been working on a network model that is simple, yet models details such as the wealth distribution in society, different classes of service, peak and off-peak traffic, and call blocking due to budgetary constraints. Parris, Keshav and Ferrari have performed experiments to study the effect of setup, per-packet and peak-load prices on the blocking probability of two classes of channels passing through a single node enforcing admission control. They have found that peak-load pricing spreads network utilization more evenly, raising revenue while simultaneously reducing the call blocking probability.

Some Tenet researchers have begun to build the foundations for the design of a new version of the Tenet real-time protocol suite. The current Tenet protocol suite implements a simple real-time communication abstraction: the simplex unicast channel. However, a large set of real-time communication applications require that the network support multicast communication, i.e., communication from a single source to multiple destinations. Examples of multicast applications are teleconferencing, distributed classrooms, and tele-shopping. To support these services, the abstraction must be extended to include multicast operation. Several challenging problems, however, have to be attacked and solved before the simplex multicast channel can be taken as the basic abstraction of the new suite.

Colin Parris and Hui Zhang have designed a Dynamic Channel Management (DCM) scheme that will be used to increase the availability and flexibility of a real-time network while maintaining the integrity of all performance contracts made to the clients with real-time channels. One method for achieving this goal requires that the network be able dynamically to modify the performance parameters and traffic characteristics of a real-time channel as well as the route traversed by the channel. These modifications will be done by the DCM scheme in a manner that is non-disruptive to the client. In most cases the modifications will be transparent to the client, unless the network deems it necessary for the client to know. This need to know is determined on the basis of the DCM policy that is in effect. DCM consists of the DCM scheme and the DCM policies. The DCM scheme is the collection of algorithms that permits the network dynamically

to modify channel parameters and routes. The modification of a channel follows a procedure whereby a real-time channel with the new performance parameters (referred to as the alternate channel) is established, the client's traffic is moved from the current real-time channel (referred to as the primary channel) to the alternate channel, and the primary channel is removed. A DCM policy is the collection of rules that determine whether a real-time channel is to be modified and the selected parameters that are to be modified. These rules may examine the network or client state data to determine whether the modification should take place. DCM policies are entirely dependent on the application, and are independent of the management algorithms. The DCM scheme consists of essentially two algorithms: (1) a channel administration algorithm, which determines whether a real-time channel can be accepted along a specified alternate route and, if so, reserves resources along the alternate route so that all of the client's traffic and performance requirements are met; the algorithm also ensures that on links that are common between the primary and alternate routes there is no duplication of resources reserved; (2) a channel transition algorithm, which ensures that the transition from the primary to the alternate channel does not violate the client's performance contract; this is especially important in the case of the delay and delay jitter bounds guaranteed in that contract. A third algorithm that is part of the DCM scheme, in a supportive rather than essential capacity, is a real-time channel routing algorithm that determines the route that would most probably satisfy the client's traffic and performance requirements. This routing algorithm considers the resources currently reserved for the primary channel when determining the alternate route [110]. These algorithms are being considered for inclusion into the next Tenet scheme.

Jean Ramaekers and Giorgio Ventre have developed a model of client-network interactions for the second version of the Tenet real-time protocol suite. The model includes new mechanisms for the establishment and runtime management of real-time connections; it is expected to reduce the complexity and the time required to establish a real-time connection, and to increase network utilization. New classes of real-time communication clients with flexible quality-of-service requirements have been introduced as a basis for the new client-network interaction model. In [111] and [112], Ramaekers and Ventre presented a new mechanism for the establishment of real-time connections in a quality-of-service network developed for the Tenet real-time protocol suite. The solution proposed is based on three ideas. The first is to increase the amount of data that a client obtains from the network when the request for a real-time channel is rejected. The second is an improved procedure for channel establishment. The third idea is the introduction of more flexible types of real-time communication service. Domenico Ferrari, Jean Ramaekers, and Giorgio Ventre are now designing a new scheme for the dynamic modification of the real-time channel characteristics. In addition to improving the utilization of network resources, this scheme should represent a first step toward the implementation of efficient mechanisms for load balancing and fault tolerance.

Eckhardt Holz, Edward Knightly, and Giorgio Ventre have designed and implemented an object oriented simulator for real-time communication networks and protocols. This simulator has two main features that distinguish it from the vast number of existing network simulators. The first is the ability to do quality-of-service simulation with an emphasis on multimedia applications. The main goal of the real-time network simulator is to show the effect of network management decisions on the traffic to verify

both the correctness and the efficiency of resource allocation schemes. In addition, work is in progress to integrate the simulator with multimedia devices and applications. For example, the simulator supports the sending of actual video packets across the simulated network. The simulator will show a “before” and an “after” picture of a video frame (or sequence of frames) in the network to facilitate the qualitative evaluation of protocols. Second, the simulator is object-oriented. This feature effectively exploits the modular and hierarchical nature of communication network protocols, and facilitates their rapid prototyping and testing.

Newton Faller is investigating how real-time operating systems similar to UNIX could provide a suitable environment for continuous-media applications. In particular he measured the latency of real-time Unix-like operating systems to better understand their capability to offer real-time communication services for multimedia applications. Continuous-media applications, when executed either locally or over a network, require an environment which should provide low delay and, sometimes more important, predictability. Computer hosts are controlled by operating systems that, supposedly, have to respond to external events with low delay and predictability. This means that such operating systems must have real-time characteristics in their ability to respond to the critical external events generated by continuous-media applications. Since the Unix operating system seems to be ubiquitous in workstations connected to networks, it is reasonable to think that its role in supporting continuous-media applications will be prevalent. However, since Unix does not exhibit real-time characteristics, it has to be modified or redesigned. Newton Faller, who has previous experience in the design and implementation of multiprocessing real-time Unix-like operating systems, has been investigating the transformations such operating systems must undergo to become appropriate to support real-time applications. He has devised a methodology [113] to evaluate how well Unix-like operating systems fit in the real-time category by measuring their latency characteristics. Special attention was given to the influence of preemption and DMA operations in interrupt acknowledge and context switching latency times. This methodology can easily be applied to any Unix-like operating system, and the way the measured data are displayed in graphical form allows an immediate comparison among systems subjected to similar workloads.

Multimedia Support: An important class of applications with real-time transport requirements needs transmission of data units at regular intervals. These applications, which we call continuous-media (CM) clients, include video conferencing, voice communication, and high-quality digital sound. Bernd Wolfinger and Mark Moran designed a data transport service for CM clients (CMTS) and its underlying protocol (CMTP). The service makes use, in particular, of an a priori characterization of future data transmission requests by CM clients. Starting from a few examples of CM clients and their specific data transmission needs, Wolfinger and Moran extracted a generalized list of data transport requirements for CM and described the basic features of a service designed to meet these requirements. CMTS provides unreliable, in-sequence, simplex, and periodic transfer of so-called stream data units (STDUs) between a sending and a receiving client, with performance guarantees on loss, delay, and throughput. An important feature of their solution is the use of shared buffers to eliminate most direct client/service interactions and to smooth traffic patterns, which may be bursty due to

fluctuations in the arrival process of data and variability of network delays [114]. A prototype of CMTP, completed by Mark Moran and Amit Gupta, is currently operational on the local Tenet testbed, although it has yet to be tested in concert with the other real-time protocols. Moran and Gupta have written a paper with Bernd Wolfinger and Domenico Ferrari that briefly describes CMTS, the interesting problems addressed in its prototype implementation, and simulation results that verify some of the gains in efficiency over message-based transport services for delivering continuous-media data [114]. Eckhardt Holz carried out a simulation study of CMTS to investigate the influence of configuration parameters on its performance, and developed a formal description of CMTP using the formal specification language SDL.

Colin Parris, Hui Zhang, and Giorgio Ventre are studying the communication needs of a client browsing high density, lossless, still images. In this application, the client cannot tolerate any loss of frames, although some marginal compression may be permitted. One of the many situations in which this occurs is in meteorological studies, where scientists wish to examine the changes in cloud formations on large satellite images, and must view all of the frames in order to observe even the minutest changes. It is assumed that the client may cause an increase by a factor K in the throughput it receives from the network when it requests an enhancement to the quality of service of its real-time channel. This request is received by the DCM algorithms, which determine whether there are sufficient resources in the network to support the request. If so, an alternate route is set up (with the higher throughput bound but the same delay and loss parameters), and a transition is made from the primary to this alternate route. The new enhanced-QOS real-time channel can then be maintained for the remaining lifetime of the channel or for a client/network-specified duration. This increase in throughput will facilitate fast-forwarding or fast-rewinding of still images. A control channel from the receiver to the sender is implemented as a real-time channel whose performance parameters provide response times of the same order as those experienced in the use of a commercial video recorder.

Mark Moran is also investigating the support required for multi-party, interactive multimedia (MIM) communication between distributed parties (e.g., computer supported collaborative work or CSCW). This work has involved gaining experience with a commercial desktop conferencing system (DECspin), and experimenting with a platform for development of multimedia applications (XMedia for DECstations). From this experience and an extensive literature search, he has characterized the requirements of canonical MIM applications and has developed a coherent design strategy for providing system and network support for MIM communication. Areas of support include time-based media, dynamic access to continuous-media I/O devices and compression modules, and real-time, multicasting communication services. This work has been done in collaboration with Riccardo Gusella at Hewlett Packard Laboratories, and is described in [115].

Flow and Congestion Control: Queueing delays have been identified by many researchers as one of the main hindrances in establishing continuous-media data connections in a rate controlled network. Anindo Banerjee and Srinivasan Keshav, a former member of the Tenet Group, now a Member of the Technical Staff at AT&T Bell Laboratories, have investigated the problem of finding the worst case end-to-end delay and

buffer occupancy bounds in networks of rate-controlled, non-work conserving servers [116]. The calculations have been based on a simple fluid model, but care was taken to make the computed delay and buffer occupancy values to be upper bounds on actual values. A simple algorithm has been devised to perform these calculations in linear time.

Multiplexing: For performance and economic reasons, ATM (Asynchronous Transmission Mode) networks must efficiently support the Internet family of protocols. Ramón Cáceres has investigated this subject by comparing the transmission efficiency achieved by a range of ATM-related protocols when transporting TCP and UDP wide-area traffic, and calculating the efficiency effects of several non-standard compression techniques [117]. The problem of how to implement these techniques cheaply is an important one. To assure an accurate workload characterization, Cáceres has driven the calculations with millions of wide-area packet lengths measured on the current Internet. He has found that networks using standard ATM procedures are inefficient in carrying wide-area data traffic: depending on the protocols used, efficiency as seen by an application program ranges between 40 and 53 percent. Moreover, due to interaction between IP datagram lengths and ATM cell padding, efficiency responds abruptly to changes in certain protocol parameters: a 4-byte increase in ATM cell payload size, for example, can yield a 10 percent increase in efficiency. Using one compression technique in isolation, the efficiency can be improved by 12 percent; simultaneously using three techniques, the improvement can be 34 percent. These issues should be taken into consideration when designing future ATM networks.

In a related project, Cáceres analyzed traces of wide-area network traffic generated by several stub networks on the current Internet. He formed a model of such traffic suitable for driving simulations of congestion control, routing, and other resource management algorithms useful in the design of current and future networks [118]. His model characterizes individual TCP conversations by the distributions of number of bytes transferred, duration, number of packets transferred, packet size, and packet interarrival time. On the basis of an analysis of wide-area network traffic traces, Cáceres finalized the workload model for driving his simulations.

Cáceres also evaluated different policies for multiplexing application-level datagram streams, or conversations, at the entrance to wide-area virtual circuit networks. Datagrams are attractive for end-to-end internetwork applications, while virtual circuits hold many advantages for the provider of integrated-services wide-area networks. When a datagram arrives at the entrance to a wide-area network, it must be routed onto an appropriate virtual circuit. It is important to provide each conversation with good performance while making efficient use of network resources. This multiplexing problem has arisen in the past and reappears with the advent of Asynchronous Transfer Mode (ATM) networks. Cáceres studied how to map a set of datagram streams onto a possibly smaller set of virtual circuits and how to choose the queueing discipline for scheduling datagrams onto these virtual circuits. His results show that the network should separate traffic types by establishing at least one virtual circuit for each type of traffic flowing between two network points of presence; conversations of the same type can share a virtual circuit. In addition, the network should use a round-robin queueing discipline to service conversations sharing a virtual circuit, and virtual circuits sharing

a communication line. This policy combines the performance advantages of allocating dedicated virtual circuits to each conversation and the resource advantages of sharing fewer virtual circuits [119].

Traffic Characterization and Generation: This project was completed at the end of 1990. There were no activities except those already mentioned in the preceding sections.

2 Project status compared to the original research plan

In the AI/Applications Group all of the projects are meeting their goals. The knowledge representation work is becoming more tightly integrated with the project on language learning and speech recognition.

The progress of the research performed by the Tenet Group has been very satisfactory, in spite of the difficulties encountered by some of the projects in their experimental phase. Indeed, as was anticipated in the previous Annual Report, the delays in the delivery of the equipment for the BLANCA testbed have hampered the testing of the Tenet real-time protocol suite. The new project on multimedia support, which has replaced the traffic characterization project, has already achieved promising results.

All research activities in the Realization group are on schedule and many major goals of the research plan have already been achieved, such as the design of the computational structures for the SPERT. Minor shifts in research emphasis will be undertaken as the need arises.

3 Prospects to reach original research goals

The prospects are good for all projects. There is growing interest in our tools and in the language work; we are monitoring this for its possible impact on our efforts.

The Tenet Group is expected to reach all of its most important research goals, even though this could require longer times than originally expected, due to the delays with which testbed equipment is being delivered and installed.

Based on the past progress and the results already achieved, the prospects to reach all envisioned research goals are excellent.

4 Impact of third party R&D results on the project

A growing interest for the Tenet real-time protocol suite has been shown by some vendors of computing and networking equipment. For example, Digital Equipment Corporation engineers are interfacing the DECspin multimedia conferencing software system to RMTP/RTIP. The collaborations activated with the vendors participating in the two major projects in which the Tenet Group is involved are expected to contribute to an increase in such interest.

The only third party R&D result that has affected our plans in any significant way is the development of a new kind of DRAM with a fast packet-based interface, called RDRAM. We will incorporate this in our designs in the coming years.

5 Changes and modifications in the research goals

Based on our past progress and goals already achieved, significant changes to the research plan are not necessary.

6 Inventions, Patent applications, etc.

No new patents since last year; however, late last year (in addition to the robust speech analysis patent application), there was a joint filing with Belgium's Lernout & Hauspie for the methods behind connectionist context-dependent phonetic probability estimation. The new work in the CNS project has a number of interesting innovations, but no patents have been filed as of yet.

7 Bibliography

- [1] Goddard, N.H.: "The Perception of Articulated Motion: Recognizing Moving Light Displays," Ph.D. Thesis & Technical Report 405, Computer Science Department, University of Rochester, 1992.
- [2] Feldman, J.: "Robots with Common Sense?," *Artificial Intelligence and Mathematical Theory of Computation*, V. Lifshitz (ed.), Academic Press, 1991.
- [3] Omohundro, S.: "Best-First Model Merging for Dynamic Learning and Recognition," *Advances in Neural Information Processing Systems 4*, Moody, J. E., Hanson, S. J., and Lippmann, R. P., (eds.), Morgan Kaufmann Publishers, San Mateo, CA, 1992, p. 958.
- [4] Uras, C. and Verri, A.: "Describing Shape through Size Functions," *Proceedings of NATO workshop on Shape in Picture*, Driebergen (NL). In press.
- [5] Regier, T.: "The Acquisition of Lexical Semantics for Spatial Terms: A Connectionist Model of Perceptual Categorization," Ph.D. thesis, Computer Science Division, EECS Department, University of California at Berkeley; also ICSI Technical Report TR-92-062.
- [6] Stolcke, A. and Wu, D.: "Tree Matching with Recursive Distributed Representations," ICSI Technical Report TR-92-025; also in working notes of the workshop on *Integrating Neural and Symbolic Processes*, National Conference on Artificial Intelligence, San Jose, Ca., July 1992.
- [7] Stolcke, A. and Omohundro, S.: "Hidden Markov Model Induction by Bayesian Model Merging," *1992 Neural Information Processing Systems 5, Natural and Synthetic*.
- [8] Omohundro, S. and Lim, C.C.: "The Sather Language and Libraries," ICSI Technical Report TR-92-017, March 1992.
- [9] Omohundro, S.: "Sather and Scientific Computation" to be published in *Computers in Physics*, September 1992.
- [10] Feldman, J.A., Lim, C.C., and Rauber, T.: "The Shared-Memory Language pSather on a Distributed-Memory Multiprocessor," Second Workshop on Languages, Compilers and Run-Time Environment for Distributed Memory Multiprocessors, September 30—October 2, 1992, Boulder, CO.
- [11] Gomes, B. and Schmidt, H.W.: "ICSIM: An Object-Oriented Connectionist Simulator," ICSI Technical Report TR-91-048, September 1991.
- [12] Bachrach, J.: "A Symbolic Complexity Analysis of Connectionist Algorithms for Distributed-Memory Machines," ICSI Technical Report TR-92-043.
- [13] Dorigo, M.: "Using Transputers to Increase Speed and Flexibility of Genetics-based Machine Learning Systems," *Microprocessing and Microprogramming*, 1992:34, 147-152, North Holland.

- [14] Bertoni, A. and Dorigo, M.: "Implicit Parallelism in Genetic Algorithms," Technical Report No. 92-012, Politecnico di Milano, Italy; Conditionally accepted for publication in the *Artificial Intelligence Journal*, Revised version submitted November, 1992.
- [15] Colorni, A., Dorigo, M. and Maniezzo, V.: "Algodesk: an experimental comparison of eight evolutionary heuristics applied to the QAP problem," Technical Report No.92-052, Politecnico di Milano, Italy; Conditionally accepted for publication in the *European Journal of Operation Research*, Revised version in preparation.
- [16] Dorigo, M., Maniezzo, V., and Montanari, D.: "Classifier-Based Robot Control Systems," 1992, to appear in *Proceedings of ICNNGA - International Conference on Neural Networks and Genetic Algorithms*, to be held in Innsbruck, Austria, 11-16 April, 1993.
- [17] Colombetti, M. and Dorigo, M.: "Learning to Control an Autonomous Robot by Distributed Genetic Algorithms", in *Proceedings of From Animals to Animats*, Second International Conference on Simulation of Adaptive Behavior (SAB92), Honolulu, HI.
- [18] Colorni, A., Dorigo, M., and Maniezzo, V.: "An Investigation of some Properties of an Ant Algorithm," in *Proceedings of the Parallel Problem Solving from Nature Conference (PPSN 92)*, Brussels, Belgium, Elsevier Publishing.
- [19] Dorigo, M. and Maniezzo, V.: "Parallel Genetic Algorithms: Introduction and Overview of Current Research," to appear in *Parallel Genetic Algorithms: Theory and Applications*, 1992, J. Stenders (Ed.), IOS Press, Amsterdam (Also Technical Report No. 92-015, Politecnico di Milano, Italy).
- [20] Dorigo, M.: "Optimization, Learning and Natural Algorithms," Ph.D. Thesis, 1992, Politecnico di Milano, Italy (in Italian).
- [21] Colombetti, M. and Dorigo, M.: "Robot Shaping: Developing Situated Agents through Learning," ICSI Technical Report No. TR-92-040; also submitted to *Artificial Intelligence Journal*, August 1992.
- [22] Dorigo, M.: "Alecsys and the AutonoMouse: Learning to Control a Real Robot by Distributed Classifier Systems," Technical Report No. 92-011, Politecnico di Milano, Italy; also submitted to *Machine Learning*, May 1992.
- [23] Dorigo, M.: "Gli Algoritmi Genetici, i Sistemi a Classificatori e il Problema dell'Animat," submitted to *Sistemi Intelligenti*, September 1992, in Italian.
- [24] Colorni, A., Dorigo, M., and Maniezzo, V.: "Introduzione agli Algoritmi Naturali," Technical Report No. 92-058, Politecnico di Milano, Italy; also submitted to *Rivista di Informatica*, September 1992, in Italian.
- [25] Colorni A., Dorigo, M., Maniezzo, V., and Trubian, M.: "Ant system for Job-shop Scheduling," submitted to *Belgian Journal of Operational Research*.

- [26] Feldman, J.A.: "The International Computer Science Institute," *Encyclopedia of Artificial Intelligence*, Second Edition, John Wiley & Sons, Inc., 1992.
- [27] Feldman, J.A.: "Structured Connectionist Models and Language Learning," in *Selected Readings of the Swedish Conference on Connectionism*, Ellis Horwood Ltd. Publishers, 1992; and contributed to *AI Review: Special Issue on the Swedish Conference on Connectionism*, 1992.
- [28] Fritzke, B.: "Kohonen feature maps and growing cell structures – a performance comparison," *Advances in Neural Information Processing 5*, 1993, Giles, Hanson, and Cowan (eds.), San Mateo, CA: Morgan Kaufmann Publishers; also presented at the Neural Information Processing Conference in Denver, December 1992
- [29] Omohundro, S. and Lim, C.C.: "The Sather Language and Libraries," ICSI Technical Report TR-92-017, March 1992.
- [30] Murer, S. and Faerber, P.: "A scalable distributed shared memory system," *Proceedings of CONPAR'92-VAPP V*, Springer LNCS 634, Lyon, September 1992.
- [31] Murer, S.: "Uebersetzung funktionaler Sprachen fuer den ADAM-Parallelrechner," Informatik Dissertationen Nr. 39, Verlag der Fachvereine, ISBN-3 7281 1947 4, Ph.D. Thesis, (in German).
- [32] Murer, S. and Marti, R.: "The FOOL Programming Language: Integrating Single Assignment and Object-Oriented Paradigms," *Proceedings of EWPC-92*, IOS Press, Barcelona, 1992.
- [33] Schmidt, H. and Omohundro, S.: "CLOS, Eiffel, and Sather: A Comparison," *The CLOS Book*, edited by Andreas Paepcke, MIT Press, November 1992; also an ICSI Technical Report, TR-91-047.
- [34] Omohundro, S.: "Sather Provides Nonproprietary Access to Object-Oriented Programming," *Computers in Physics*, September 1992, p. 444.
- [35] Moody, J. and Utans, J.: "Principled Architecture Selection for Neural Networks: Application to Corporate Bond Rating Prediction," *Advances in Neural Information Processing Systems 4*, Moody, J.E., Hanson, S.J., and Lippmann, R.P. (eds.), Morgan Kaufmann Publishers, San Mateo, CA, 1992.
- [36] Utans, J. and Gindi, G.: "A Neural Network Approach to Object Recognition and Image Partitioning within a Resolution Hierarchy," *Proceedings of SPIE Intelligent Information Systems OR'92*, Orlando, FL, April 20-24, 1992.
- [37] Utans, J. and Gindi, G.: "A Neural Network Performs Context Guided Segmentation of Shapes via Bayesian Inference," presented at *Neural Networks for Computing*, Snowbird, Utah, April 7-10, 1992.
- [38] Utans, J., and Moody J.: "Selecting Neural Network Models for the Prediction of Corporate Bond Ratings," invited talk at the *33rd Joint National Meeting, The Institute of Management Science/Operations Research Society of America (TIMS/ORSA)*, Orlando, FL, April 26-29, 1992.

- [39] Verri, A., Straforini, M., and Torre, V.: "Computational aspects of motion perception in artificial and natural visual systems," *Phil. Trans. R. Soc. B.*, 1992:337, 429-443.
- [40] Orban, G.A., Lagae, L., Verri, A., Raiguel, S., Xiao, D., Maes, H., and Torre, V.: "First order analysis of optical flow in monkey brain," *Proc. Natl. Acad. Sci. USA*, 1992:89, 2595-2599.
- [41] Campani, M. and Verri, A.: "Motion analysis from first order properties of optical flow," *CVGIP: Image Understanding*, 1992:56, 90-107.
- [42] Uras, C. and Verri, A.: "Studying Shapes through Size Functions," *Proceedings of the NATO Workshop "Shape in Picture"*, Driebergen (The Netherlands), Sept. 7-11, 1992.
- [43] Rognone, A., Campani, M., and Verri, A.: "Identifying multiple motions from optical flow," *Proceedings of 2nd European Conference on Computer Vision*, S. Margherita (Italy), 1992.
- [44] Malisia, A., Campani, M., and Verri, A.: "A simple and robust method for estimating the motion of a mobile vehicle from optical flow," *Proceedings of 2nd Workshop on Robust Computer Vision*, Bonn, 1992.
- [45] Verri, A., Straforini, M., and Torre, V.: "A model of the spatial organisation of the receptive fields of neurons in MST," *15th European Conference on Visual Perception*, Pisa, 1992.
- [46] Culler, D., Karp, R., Patterson, D., Sahay, A., Schauser, K.E., Santos, E., Subramonian, R., von Eicken, T., "LogP: Towards a Realistic Model of Parallel Computation," to appear as a U.C. Berkeley Computer Science Division Report. Also, to be presented at ACM Symp. on Principles and Practices of Parallel Programming, May, 1993.
- [47] Karpinski, M., Dahlhaus, E., and Kelsen, P.: "An Efficient Parallel Algorithm for Computing a Maximal *Proceedings of Beijing IJCNN*, 1992, (in press).
- [48] Karp, R., Luby, M., and Meyer auf de Heide, F.: "Efficient PRAM Simulation on a Distributed Memory Machine" a preliminary version appears in 24th *STOC*, 1992. Invited to appear in a special issue of *Algorithmica* devoted to "Randomized Algorithms". Journal version in preparation.
- [49] Luby, M. and Nisan, N.: "A Parallel Approximation Algorithm for the Positive Linear Programming Problem," accepted to *STOC* 1993.
- [50] Bertoni, A., Campadelli, P., Morpurgo, A., and Panizza, S.: "Polynomial Uniform Convergence and Polynomial-Sample Learnability," *Proceedings 5th ACM Workshop on Computational Learning Theory*, pages 265-271, ACM, 1992. An extended version is in preparation as an ICSI technical report.

- [51] Bshouty, N.H., Hancock, T.R., Hellerstein, L., and Karpinski, M.: "Read Once Threshold Formulas, Justifying Assignments, and Generic Transformations," ICSI Technical Report TR-92-020; also submitted to *Computational Complexity*, 1992.
- [52] Werther, T.: "VC Dimension of Sparse Sums of Chebyshev Polynomials," first draft.
- [53] Karpinski, M., and Werther, T.: "VC Dimension and Sampling Complexity of Learning Sparse Polynomials and Rational Functions," ICSI Technical Report TR-91-044, August 1991, to appear in *SIAM Journal on Computing*.
- [54] Werther, T.: "VC Dimension and Learnability of Sparse Polynomials," Ph. D. Thesis, University of Bonn, 1991.
- [55] Even, G., Goldreich, O., Luby, M., Nisan, N. and Veličković, B.: "Approximations of General Independent Distributions." A preliminary version appears in the proceedings of the 24th *STOC*, 1992.
- [56] Linial, N., Luby, M., Saks, M., and Zuckerman, D.: "Efficient construction of a small hitting set for combinatorial rectangles in high dimension," accepted to *STOC* 1993.
- [57] Luby, M., Wigderson, A. and Veličković, B.: "Deterministic Approximate Counting of Depth-2 Circuits," submitted.
- [58] Luby, M., Naor, M. and Naor, S.: "On Removing Randomness from a Parallel Algorithm for Edge Connectivity." Work in progress, 1992.
- [59] Luby, M.: "Pseudo-randomness and Applications," a monograph to be published by *Princeton University Press*. Currently in final stages of revision, projected publication date is June 1993, approximately 200 pages.
- [60] Herzberg, A., and Luby, M.: "Public Randomness in Cryptography," presented at *CRYPTO 1992*, August, 1992, to appear in proceedings.
- [61] Rabinovich, Y., Sinclair, A., and Wigderson, A.: "Quadratic Dynamical Systems," *Proceedings of IEEE Symposium on Foundations of Computer Science*, Pittsburgh, October 1992.
- [62] Kenyon, C., Randall, D., and Sinclair, A.: "Matchings in Lattice Graphs," to appear in *Proceedings of ACM Symposium on Theory of Computing*, San Diego, May 1993.
- [63] Buergisser, P., Lickteig, T., and Karpinski, M.: "On the Randomized Algebraic Test Complexity," first draft.
- [64] Karpinski, M., and Goldmann, M.: "Fooling Small Sums and Derandomizing Majority Circuits," first draft.
- [65] Aberer, K., and Codenotti, B.: "Models of Computations in Numerical Analysis," submitted for publication.

- [66] Aberer, K., and Codenotti, B.: "Towards a Complexity Theory for Numerical Analysis," ICSI Technical Report TR-92-012.
- [67] Aberer, K.: "Combinatory Differential Fields: An Algebraic Approach to Approximate Computation and Constructive Analysis," ICSI Technical Report TR-91-061.
- [68] Aberer, K.: "Combinatory Models and Symbolic Computation," *Proceedings of DISCO 92*.
- [69] Aberer, K., and Codenotti, B.: "On the Power of Discontinuous Approximate Computations," ICSI Technical Report TR-92-026; also submitted for conference publication.
- [70] Alizadeh, F., Karp, R., Newberg, L., and Weisser, D.: "Physical Mapping of Chromosomes: A Combinatorial Problem in Molecular Biology," to appear in *Algorithmica*. Preliminary version to be presented at ACM/SIAM Symposium on Discrete Algorithms, January, 1993.
- [71] Karp, R.: "A Generalization of Binary Search," submitted to ACM/SIAM Symposium on Discrete Algorithms July, 1992.
- [72] El-Yaniv, R., Fiat, A., Karp, R., and Turpin, G.: "Competitive Analysis of Financial Games," to appear in *Proceedings of the IEEE Symposium on FOCS*, October, 1992.
- [73] Karp, R.: "On-Line Algorithms Versus Off-Line Algorithms: How Much is It Worth to Know the Future," to appear in *Proceedings of World Computer Congress*, September, 1992.
- [74] d'Amore, F., Marchetti-Spaccamela, A., Nanni U.: "The Weighted List Update Problem and the Lazy Adversary", to appear in *Theoretical Computer Science*. ICSI Technical Report TR-92-011. A preliminary version of this paper appears in *Proceedings of the Workshop on Algorithms and Data Structures (WADS 91)*, Carleton University, Ottawa, Canada, 14-16 August 1991; *Lecture Notes in Computer Science*, 519, Springer Verlag, Berlin, 1991.
- [75] Marchetti-Spaccamela A., Nanni U., and Rohnert H.: "On-line Graph Algorithms for Incremental Compilation", ICSI Technical Report TR-92-056.
- [76] Karpinski, M., Grigoriev, D., and Singer, M.: "Computational Complexity of Sparse Algebraic Function Interpolation," *Proceedings of MEGA'92*, May 1992.
- [77] Karpinski, M., Grigoriev, D., and Odlyzko, A.M.: "Existence of Short Proofs for Nondivisibility of Sparse Polynomials under the ERH," *Proceedings of ISSAC'92*, pp. 117-122, July 1992.
- [78] Ausiello, G., Italiano, G. F., Marchetti-Spaccamela, A., and Nanni, U.: "Incremental Algorithms for Minimal Length Paths", *Journal of Algorithms*, vol. 12, no. 4 (December 1991), 615-638.

- [79] Ausiello, G., Italiano, G. F., and Nanni, U.: "Optimal Traversal of Directed Hypergraphs" ICSI Technical Report TR-92-073, 1992.
- [80] Alimonti, P., Feuerstein, E., Nanni, U.: "Linear Time Algorithms for Liveness and Boundedness in Conflict-Free Petri Nets", proceedings *1st Latin American Theoretical Informatics (LATIN 92)*, São Paulo, Brazil, April 6-10, 1992. *Lecture Notes in Computer Science*, 583, Springer Verlag, Berlin, 1992. ICSI Technical Report TR-92-008, 1992.
- [81] Morgan, N., Beck, J., Kohn, P., Bilmes, J., Allman, E. and Beer, J.: "The Ring Array Processor (RAP): A Multiprocessing Peripheral for Connectionist Applications", *Journal of Parallel and Distributed Computing*, Special Issue on Neural Networks, 1992, v14, pp. 248-259, 1992.
- [82] Asanovic, K., Beck, J., Kingsbury, B., Kohn, P., Morgan, N. and Wawrzynek, J.: "SPERT: A VLIW/SIMD Neuro-Microprocessor", *Proc. of IJCNN '92*, II-577-582.
- [83] Asanovic, K., Beck, J., Kingsbury, B., Kohn, P., Morgan, N. and Wawrzynek, J.: "SPERT: A VLIW/SIMD Microprocessor for Artificial Neural Network Computations", *Proc. of Intl. Conf. on Application Specific Array Processors*, pp. 178-190, Berkeley, 1992.
- [84] Asanović, K., Morgan, N. and Wawrzynek, J.: "Using Simulations of Reduced Precision Arithmetic to Design a Neuro-Microprocessor", *Journal of VLSI Signal Processing*, 1993, to appear.
- [85] Wawrzynek, J., Asanovic, K. and Morgan, N.: "The Design of a Neuro-Microprocessor", *IEEE Trans on Neural Networks*, in press.
- [86] Afghahi, M. and Svensson, C.: "A Unified Single-Phase Clocking Scheme for VLSI Systems", *IEEE JSSC*, 25(1):225-233, February 1990.
- [87] Hermansky, H., Morgan, N., Bayya, A. and Kohn, P.: "RASTA-PLP Speech Analysis Technique", *IEEE Intl. Conf. on Acoustics, Speech, & Signal Processing*, San Francisco, 1992, I-121-124.
- [88] Hermansky, H. and Morgan, N.: "Towards Handling the Acoustic Environment in Spoken Language Processing", *International Conference on Spoken Language Processing*, pp. 85-88, 1992.
- [89] Hermansky, H., Hirsch, H. and Morgan, N.: "Recognition of Speech in Additive and Convolutional Noise Based on RASTA Spectral Processing", *IEEE Intl. Conf. on Acoustics, Speech, & Signal Processing*, 1993, in press.
- [90] Hermansky, H. and Morgan, N.: "RASTA Extensions: Robustness to Additive and Convolutional Noise", *Proceedings of Workshop on Speech Processing in Adverse Conditions*, pp. 115-118, 1992.
- [91] Trompf, M.: "Experiments with Noise Reduction Neural Networks for Robust Speech Recognition", ICSI Technical Report TR-92-035.

- [92] Trompf, M.: "Neural Network Development for Noise Reduction in Robust Speech Recognition", *Proc. of IJCNN '92*, IV-722-727.
- [93] Bourlard, H., Morgan, N. and Renals, S.: "Neural Nets and hidden Markov models: Review and generalizations", *Speech Communication*, vol. 11, no. 2-3, pp. 237-246.
- [94] Renals, S., Morgan, N., Cohen, M., Franco, H. and Bourlard, H.: "Connectionist Probability Estimation in the DECIPHER Speech Recognition System", *IEEE Intl. Conf. on Acoustics, Speech, & Signal Processing*, San Francisco, 1992, I-601-604.
- [95] Renals, S., Morgan, N., Cohen, M., Franco, H. and Bourlard, H.: "Improving Statistical Speech Recognition", *Proc. of IJCNN '92*, II-302-307.
- [96] Bourlard, H., Morgan, N., Wooters, C. and Renals, S.: "CDNN: A Context Dependent Neural Network for Continuous Speech Recognition", *IEEE Intl. Conf. on Acoustics, Speech, & Signal Processing*, San Francisco, 1992, II-349-352.
- [97] Franco, H., Cohen, M., Morgan, N., Renals, S. and Abrash, V.: "Context-Dependent Connectionist Probability Estimation in a Hybrid HMM-Neural Net Speech Recognition System", *Proc. of Beijing IJCNN*, 1992, in press.
- [98] Morgan, N. and Bourlard, H.: "Factoring Networks by a Statistical Method", *Neural Computation*, vol. 4, no. 6, pp. 835-838, 1992.
- [99] Cohen, M., Franco, H., Morgan, N., Rumelhart, D. and Abrash, V.: "Hybrid Neural Network/Hidden Markov Model Continuous Speech Recognition", *Proc. of Intl. Conf. on Speech and Language Processing*, Banff, Canada, pp. 915-918, 1992.
- [100] Konig, Y. and Morgan, N.: "GDNN: A Gender-Dependent Neural Network for Continuous Speech Recognition", *Proc. of IJCNN '92*, II-332-337.
- [101] Konig, Y., Morgan, N., Wooters, C., Abrash, V., Cohen, M. and Franco, H.: "Modeling Consistency in a Speaker Independent Continuous Speech Recognition System", *Advances in Neural Information Processing Systems 5*, in press.
- [102] Wooters, C. and Morgan, N.: "Connectionist-Based Acoustic Word Models", *IEEE Workshop on Neural Networks for Signal Processing*, pp. 157-163, Copenhagen, 1992.
- [103] Wooters, C. and Morgan, N.: "Acoustic Sub-word Models in the Berkeley Restaurant Project", *International Conference on Spoken Language Processing*, pp. 1551-1554, 1992.
- [104] Ferrari, D., Pasquale, J. and Polyzos, G.: "Network Issues for Sequoia 2000," *Proceedings of COMPCOM 92*, pages 401-406, 1992.
- [105] Ferrari, D.: "Real-Time Communication in an Internetwork," *Journal of High Speed Networks*, no. 1, vol. 1, pages 79-103, 1992.

- [106] Ferrari, D.: "Design and Application of a Delay Jitter Control Scheme for Packet Switching Internetworks," *Proceedings of the Second International Workshop on Network and Operating System Support for Digital Audio and Video*, Heidelberg, Germany, pages 72–83, November 1991.
- [107] Ferrari, D. and Zhang, H.: "The Rate Controlled Static Priority Discipline," ICSI Technical Report TR-92-003, January 1992.
- [108] Zhang, H. and Fisher, T.: "Preliminary Measurement of the Real-Time Internet Protocol," *Proceedings of the Third International Workshop on Network and Operating System Support for Digital Audio and Video*, San Diego, CA, November 1992.
- [109] Parris, C., Keshav, S. and Ferrari, D.: "A Framework for the Study of Pricing in Integrated Networks," ICSI Technical Report TR-92-016, March 1992.
- [110] Parris, C., Zhang, H. and Ferrari, D.: "A Mechanism for Dynamic Re-routing of Real-time Channels," ICSI Technical Report TR-92-053, August 1992.
- [111] Ramaekers, J. and Ventre, G.: "Quality-of-Service Negotiation in a Real-Time Communication Network," ICSI Technical Report TR-92-023, April 1992.
- [112] Ramaekers, J. and Ventre, G.: "Client-Network Interaction in a Real-Time Communication Environment," *Proceedings of IEEE GLOBECOMM '92*, 1992.
- [113] Faller, N.: "Measuring the Latency Time of Real-Time Unix-like Operating Systems", ICSI Technical Report TR-92-037, June 1992.
- [114] Ferrari, D., Gupta, A., Moran, M. and Wolfinger, B.: "A Continuous Media Communication Service and its Implementation," *Proceedings of IEEE GLOBECOM '92*, 1992.
- [115] Moran, M. and Gusella, R.: "System Support for Efficient Dynamically-Configurable Multi-Party Interactive Multimedia Applications," *Proceedings of the Third International Workshop on Network and Operating Systems Support for Digital Audio and Video*, San Diego, November 1992.
- [116] Banerjea, A. and Keshav, S.: "Queuing Delays in Rate Controlled Networks", ICSI Technical Report TR-92-015, February 1992.
- [117] Caceres, R.: "Efficiency of ATM Networks in Trasporting Wide-Area Traffic", ICSI Technical Report TR-91-043, July 1991.
- [118] Danzig, P., Jamin, S., Caceres, R., Mitzel, D. and Estrin, D.: "An Empirical Workload Model for Driving Wide-Area TCP/IP Network Simulations", *Journal of Internetworking Research and Experience*, 1992.
- [119] Cáceres, R.: "Multiplexing Traffic at the Entrance to Wide-Area Networks", Ph.D. Dissertation, University of California, Berkeley, October 1992.

- [120] Banerjea, A. and Mah, B.A.: “The Design of a Real-Time Channel Administration Protocol”, *Proceedings of Second International Workshop on Network and Operating System Support for Digital Audio and Video*, Heidelberg, November 1991.