First Annual Progress Report

ALCOM-FT

Algorithms and Complexity
Future Technologies

Project No. IST-1999-14186

August 2001

Summary

This is the first annual progress report for the ALCOM-FT project, supported by the European Commission as project number IST-1999-14186 under the Information Society Technologies programme of the Fifth Framework. The report covers the period June 1, 2000 to May 31, 2001.

ALCOM-FT brings together ten of the leading groups in algorithms research in Europe in a project that proposes to discover new algorithmic concepts, identify key algorithmic problems in important applications, and contribute to the accelerated transfer of advanced algorithmic techniques into commercial systems. The ten participants of the project are listed in Table 1.

No.	*Full Name*	*Short Name*
1	BRICS, Department of Computer Science, University of Aarhus (coordinating site)	Aarhus
2	Department of Software, Polytechnic University of Catalunya, Barcelona	Barcelona
3	Department of Computer Science, University of Cologne	Cologne
4	INRIA, Rocquencourt	INRIA
5	Max-Planck-Institut für Informatik, Saarbrücken	MPI
6	Department of Mathematics and Computer Science, University of Paderborn	Paderborn
7	Computer Technology Institute, Patras, Greece.	CTI
8	Department of Computer and System Science, University of Rome "La Sapienza"	Rome
9	Department of Computer Science, University of Utrecht	Utrecht
10	Department of Computer Science, University of Warwick	Warwick

Table 1: The participants of the ALCOM-FT project.

During the first twelve months of the project, the work has been carried out as planned. Proper coordination of the project has been ensured through the election of a Consortium Board, a Steering Committee and a Work Package Leader for each work package of the project. Dissemination of the work done has been ensured by the publication of scientific reports, and by the creation of websites for the entire project as well as for several tasks within the project.

The twelve deliverables relating to the first year of the project are listed in Table 2. Of these, all except the pair D4 and D10 have been delivered within the month promised. The deliverables D4 and D10 involve the publication of a volume in the *Lecture Notes in Computer Science* series of *Springer Verlag*, and have been delayed due to the time involved in the editorial process. The deliverable D1 represents 174 scientific reports, all of which are available online.

The rest of this document is organized into three parts. The first describes the progress achieved within the various work packages of the project. The second list significant events during the year for each participating site. The third list the scientific reports published as part of the project.

No.	*Deliverable*	*Month*
D1	Research reports	Cont.
D2	Project Presentation	3
D3	Dissemination and Use Plan	6
D4	Guidelines for algorithmic experiments (internal release)	6
D5	Testbed for experimental algorithmics (specification)	6
D6	Ext. memory experimental platform (design)	12
D7	Distributed algorithmic engineering software package (prototype)	12
D8	BSP-style library for dynamic distributed environments (prototype)	12
D9	Production and transportation planning modeling report	12
D10	Guidelines for algorithmic experiments (external release)	12
D11	Algorithm Forum web site	12
D12	First progress report	12

Table 2: The deliverables relating to the first year.

Contents

Part I Work Packages

WP 1: Massive Data Sets

Participants: Aarhus, Barcelona, INRIA, MPI, Paderborn, CTI, Rome, Warwick
Coordinating site: Barcelona
Work Package Leader: José L. Balcazar
Number of technical reports: 28, of which 21 also appear in other work packages

The work performed corresponds well to the planned efforts of this work package as described by the contract. Each of the tasks specified has been successfully addressed and led to the production of several reports and, in many cases, software. Work has advanced quite fast, and in some cases has received wide recognition, in algorithms for parallel disks, in graph problems, and, to a lesser extent, in extraction of relevant information from unstructured or poorly structured massive data; we will give somewhat more detailed accounts of these parts of the work, which correspond to roughly half the reports of WP 1, and only very briefly mention, later on, the work done on the remaining tasks. It is expected that these other tasks will be treated in more depth in future yearly reports, as their advance becomes more prominent. A specific section describes the status of the software deliverable LEDA-SM implementing secondary storage data structures.

Many of the reports associated with this work package describe work including a part of experimental implementation, and there is an equilibrated split into proper advance of this work package and healthy cooperation with other work packages, as can be seen by the number of reports related to multiple work packages.

Algorithms for Parallel Disks

One focus of our research has been algorithmic support for parallel disks. Our approach is to use randomization, redundant storage, and sophisticated scheduling algorithms to achieve automatic load balancing. It turns out that an almost trivial algorithm is optimal for writing a sequence of L blocks to D disks [79]. For random placement this translates into an algorithm working in $(1 + O(D/m))L/D$ output steps if m buffer frames are available and $L = \omega(D \log D)$. Using a newly discovered duality between writing to disks and prefetching, we get an optimal off-line prefetching algorithm with the same efficiency. Using additional refinements for allocating data streams we obtain the best currently known parallel disk sorting algorithms. Very efficient read accesses to parallel disks can be achieved even without knowledge of future accesses if two randomly placed copies of each block are available. We explain how to do this with efficiency close to one, even for variable length blocks and disks connected to the processors using a hierarchical communication network [82]. The system retains most of its efficiency when a system component fails and supports fast failure recovery without replacing the failed component. The above analytical results have been obtained for batched access to the disks. For continuous streams of requests we have developed corresponding algorithms [99] that are analyzed using extensive simulations. All the algorithms can be generalized to allocation strategies with lower redundancy and higher fault tolerance. The variant for the asynchronous case may be of some independent theoretical interest since it optimally solves a generalized form of the maximum cardinality bipartite matching problem in polynomial time.

Graph Algorithms

We consider the classical graph traversal problems breadth first search (BFS), depth first search (DFS) and single-source shortest-paths (SSSP) for large inputs. Parallel computing may help to solve the problems faster and/or provide the aggregated memory in order to avoid slow external memory computing on hard disks. In case external memory access is unavoidable, we face a challenging algorithmic problem even for BFS and DFS. Whereas BFS and DFS are easily solved sequentially using linear time, fast and efficient parallel and external graph exploration still constitutes one of the big algorithmic challenges. Even for purely sequential SSSP on directed graphs with n nodes, m edges and non-negative edge weights, all known algorithms require super-linear time in the worst-case.

In [98] we provide the first SSSP algorithm that provably achieves linear $O(n + m)$ average-case time on arbitrary directed graphs with random edge weights. The label-correcting algorithm applies a bucket-based approximate priority queue with adaptive bucket-splitting: nodes with higher degree are treated in buckets of smaller width, thus simultaneously limiting the overhead for node re-insertions and bucket traversals. A parallel version of a precursor algorithm with fixed bucket width was presented in [95]. An improved parallel SSSP algorithm using the adaptive step width idea of [98] is given in [94]. For power-law graph models of the WWW or call graphs this results in the first work-efficient $o(n^{1/4})$ average-case time algorithm.

In [96] we give the first external memory BFS algorithm which achieves $o(n)$ I/Os on arbitrary undirected graphs with constant maximum node degree d. Let M and $B > d$ denote the main memory size and block size, respectively. Having $\mathrm{Sort}(x) = \Theta(\frac{x}{B} \cdot \log_{M/B} \frac{x}{B})$, our algorithm needs $\mathcal{O}(\frac{n}{\gamma \cdot \log_d B} + \mathrm{Sort}(n \cdot B^\gamma))$ I/Os and $\mathcal{O}(n \cdot B^\gamma)$ external space for an arbitrary parameter $0 < \gamma \leq 1/2$. Using an algorithm of [38], the result carries over to BFS, DFS and SSSP on undirected embedded planar graphs with arbitrary node degrees. External memory DFS on undirected embedded planar graphs can be solved even more efficiently [93] by properly adapting a non-optimal parallel algorithm: based on an I/O efficient method to find simple cycle separators in biconnected planar graphs we show how to compute the DFS tree using $O(\mathrm{Sort}(n) \log n)$ I/Os and linear space. Furthermore, we provide an $O(\mathrm{Sort}(n))$ I/O reduction from DFS to BFS on undirected embedded planar graphs thus showing that for planar graphs BFS is as hard as DFS.

Poorly-Structured Large Datasets

An important problem in the handling of large data sets is that of extracting relevant information from largely unstructured data, where the treatment often has to be based on statistical decisions. One then needs to be able to recognize what are statistically unavoidable regularities and distinguish these from the observed events that are likely to be meaningful. Very general and original methods have been developed in order to quantify precisely what to expect regarding the occurrences of patterns in large texts obeying a variety of probabilistic models. This permits us to set thresholds beyond which detection of a pattern is to be considered meaningful. Applications are in the area of intrusion detection [43] where large audit files are searched for specific intrusion patterns that are interspersed with unpredictable irrelevant information; the corresponding patterns are appreciably richer than the one classically considered in that they allow constrained gaps between atomic events. Another range of application is to large bases of biological sequences where probability-based decision making is obviously essential [159]. It is to be noted that classical probabilistic methods do not readily provide answers to these questions, while the original combination of formal and analytic approaches developed in the project enables us to attack the analysis of significance thresholds for an extremely wide class of complex patterns.

On relational data, a central problem in Data Mining applications is the computation of frequent sets from a given family of subsets of a fixed universe. This is actually a version of the problem of computing the transversal hypergraph of a given hypergraph, implicitly described by the frequencies of occurrence of items from the universe in the dataset. There is already a large number of algorithms for this problem, but there is some agreement in the research community that these solutions are not completely satisfactory since they do not scale up well. We have proposed in [39] an alternative scheme based on a best-first strategy on which additional experimentation is ongoing, both on algorithmic aspects and on refining the statement of the problem in search of more efficient algorithmics. We have also studied an alternative for classification problems based on a theoretical mechanism from Computational Learning, the Support Vector Algorithms. In [92] we have provided an alternative randomized decomposition algorithm for which we could prove a strong bound on the expected completion time: it is the first decomposition algorithm for this model for which expected running time bounds are proved, the previous state-of-the-art results only achieving proofs of convergence with no guarantee on the running time. The bounds obtained suggest better scalability results for these algorithms, and we are currently developing additional theoretical work and setting up the necessary experimentation platform to practically validate our theoretical results.

Software Libraries for External Memory

We have completed the single disk external memory library LEDA-SM (Library of Efficient Data types and Algorithms — Secondary Memory) [73]. The library consists of a low level block manager that allows the use of file systems and raw devices interchangeably and a higher level part that implements numerous algorithms and data structures such as sorting, matrix arithmetics, simple graph algorithms, priority queues, B-trees, buffer trees, and and suffix arrays. For a more detailed discussion refer to Deliverable D6.

Other Work

From the reports not mentioned so far, let us simply explain that they range from highly abstract long-term research to practically motivated specific problems. On one extreme we could point out the novel algebraic approaches to Computational Learning Theory problems, both with respect to interaction models to obtain information on the target representation to be learned, and with respect to the knowledge representation mechanisms themselves; on the other, advances on the study of data structures of wide applicability and on issues of parallelization in the BSP model or in network models where load balancing and data placement are crucial. More detailed information on all this work can be obtained from the repository of technical reports on the project web page.

WP 2: Networks and Communication

Participants: INRIA, MPI, Paderborn, CTI, Rome, Warwick
Coordinating site: CTI
Work Package Leader: Christos Kaklamanis
Number of technical reports: 75, of which 47 also appear in other work packages

The principal aim of the work within WP 2 is to design, theoretically analyze, and experimentally validate efficient and robust solutions to selected fundamental optimization issues in modern network communications. The work during the first year of the projects focused on two central issues, *efficient communication in modern networks* and *dynamic distributed environments*, addressing the topics

- Wavelength optimization in all–optical networks.
- Frequency spectrum management in wireless networks.
- Bandwidth allocation in high–speed (ATM) networks.
- Stochastic modelling of networks and telecommunication protocols.
- Communication primitives.
- Static and dynamic load–balancing.
- Thread scheduling.
- Data management in networks.
- BSP algorithms.

- Prefetching and caching in the WWW.
- Distributed algorithms.

In the rest of this section, we highlight important developments on the above topics focusing on the theoretical developments.

Wavelength optimization in all–optical networks

A crucial problem in WDM (Wavelength Division Multiplexing) all–optical networks is the wavelength routing problem. In recent work [54, 59] (see also [55]), we study the path coloring problem in trees, an important variant of the wavelength routing problem. The path coloring problem can be formalized as follows: We are given a directed tree T and a set R of directed paths on T. We wish to assign colors to paths of R, in such way that no two paths that share a directed arc of T are assigned the same color, and the total number of colors used is minimized. The problem has been proved to be NP–complete even for the case of binary trees. In the bibliography, the results for the problem are expressed in terms of the maximum load l of R, i.e. the maximum number of paths that go through a directed arc of T.

So far, only deterministic greedy algorithms had been presented for the problem. The best known algorithm colors any set R of maximum load l using at most $5l/3$ colors. Alternatively, we say that this algorithm has performance (or approximation) ratio 5/3. It was also known that no deterministic greedy algorithm can achieve a performance ratio better than 5/3.

One direction of research [59] was to define the class of greedy algorithms that use randomization. The results of this research are expressed in terms of the depth of the tree and of the maximum load l of R, i.e. the maximum number of paths that go through a directed arc of T.

The limitations of randomized greedy algorithms were studied and it was proved that, with high probability, randomized greedy algorithms cannot achieve a performance ratio better than 3/2 when applied for binary trees of depth $\Omega(l)$, and $1.293 - o(1)$ when applied for binary trees of constant depth.

Exploiting inherent properties of randomized greedy algorithms, the first randomized algorithm for the problem was obtained. This algorithm uses at most $7l/5 + o(l)$ colors for coloring any set of paths of maximum load l on binary trees of depth $O(l^{1/3-\epsilon})$, with high probability. An existential upper bound of $7l/5 + o(l)$ was also presented that holds on any binary tree.

For the analysis of the above upper and lower bounds, tail inequalities for random variables following hypergeometrical probability distributions were developed.

Another direction of research [59] was to try to approximate the optimal solution for path coloring on bounded degree trees based on an optimal solution for the corresponding fractional path coloring problem. We showed how to optimally compute the fractional chromatic number of a set of paths R on a tree of bounded degree. The main idea of this technique is to inductively construct a linear program whose solution is actually the optimal solution to the fractional path coloring problem. The linear program produced is proved to have polynomial size (i.e., polynomial number of variables) as far as the degree of the tree is bounded. This technique can be generalized to graphs of bounded degree and bounded treewidth. Then, the (integral) coloring problem can be approximated within a small factor (which is 1.61 for trees and 2.22 for trees of rings) by applying the well–known randomized rounding technique to the fractional path coloring.

Frequency spectrum management in wireless networks

Frequency spectrum is a scarce resource for wireless networks utilizing frequency division multiplexing (FDM) technology. We have studied solutions for two problems related to frequency spectrum optimization in wireless networks.

The frequency allocation problem is stated as follows: Given calls (i.e., users that wish to communicate) in a cellular network, assign frequencies to all users in such way that users in the same or adjacent cells are assigned distinct frequencies and the total number of frequencies used is minimized. The on-line version of the problem was considered where call frequencies are not known in advance and the algorithm must immediately assign a frequency to a new call when it appears. For the greedy frequency allocation algorithm, we showed that its competitive ratio is 2.5 against offline adversaries.

In [108, 109], we studied the off-line version of frequency allocation problems with constraints and proved complexity results and approximation algorithms with constant for planar network topologies. The optimization problem was formalized as the following graph-theoretic problem: Given a graph whose vertices represent users that wish to communicate and whose edges indicate that two users are close to each other, assign colors (i.e., positive integers) to the vertices of the graph such that users at distance one are assigned colors with distance two between then and users at distance two are assigned different colors. The objective is to minimize the total number of colors.

In [57, 60], we designed and analyzed randomized on–line algorithms for the call control problem in cellular networks. Given calls (i.e., users that wish to communicate) in a cellular network with w frequencies, the call control problem is to find a frequency allocation for some of the users such that the number of users served is maximized. Using competitive analysis, it was proved that there exist a series of randomized algorithms that beats the deterministic lower bound on the competitiveness achieving a competitive ratio of 2.65 for cellular networks with one frequency. In the same work, several lower bounds were derived (for planar and cellular networks) while the analysis was also extended to arbitrary network topologies.

Bandwidth allocation in high–speed (ATM) networks

In [47], we study a problem motivated by the allocation of request for bandwidth slots on a shared satellite channel. Communication requests are scheduled on non-intersecting rectangles in the time/bandwidth Cartesian space. The goal is that of maximizing the benefit obtained from accepted requests. This problem turns out to be equal to the maximization version of the well known Dynamic Storage Allocation problem when storage size is limited and requests must be accommodated within a prescribed time interval.

The authors present constant approximation algorithms for the problem introduced using as a basic step the solution of a fractional Linear Programming formulation. This problem has also been independently studied with different techniques by Bar-Noy et al. (STOC 00). The new approach approach gives the best approximation ratio known for the problem.

Stochastic modelling of networks and telecommunication protocols

In this context we aim at investigating the basic features of mathematical models describing telecommunication networks. The objective is firstly to give a precise description of the impact of the basic algorithms controlling the telecommunication networks and secondly to optimize the choice of their parameters. The equilibrium properties and transient behavior of the stochastic models of networks and telecommunication protocols are analyzed through Markovian methods. The renormalization (Euler's scaling) is the main fundamental tool to investigate the behavior of complex multi-dimensional Markov processes involved.

We investigated the following problems:

a) A fundamental study of the renormalized processes involved in the stochastic models describing telecommunication networks [10, 44].

b) Allocation bandwidth problems in a TCP like network [16, 33].

We have analyzed the fluid limits of the Markov processes describing an allocation scheme and a telecommunication network. The fluid limits approach is a first order analysis of the state of complex networks. It is one of the most promising methods to tackle these difficult problems. If $Z(t)$ describes the state of the network at time t, the state and the time are scaled by the size of the initial state, $N = \|Z(0)\|$, i.e. one has to study $Z(tN)/N$ as N goes to infinity. In this manner, roughly speaking, a stochastic process is "almost" reduced to a deterministic dynamical system. A basic question concerns the relations between the respective properties of the initial process and the dynamical system.

The main results obtained in this part are the following

- An old conjecture concerning the non randomness of fluid limits in a classical Markovian model of a telecommunication network has been shown to be false [44]. Randomness implies that deterministic differential equations cannot describe completely the first order behavior of the network.

- It was shown that for some allocation schemes, the limiting dynamical system involves products of random matrices. A complete study of the initial stochastic process has been achieved with the help of the dynamical system [10].

The second part concerns specifically TCP (Transmission Control Protocol) networks. Two approaches have been used:

- A local study of TCP [33]. The protocol is analyzed in the following situation: a source is sending packets to a destination, it is assumed that the rest of the network interacts through the loss of some of the packets. Using a Markovian framework we have analyzed the throughput of this protocol when the loss probability α of a packet tends to 0. As a result we have proved that the asymptotic throughput is $\sim 1.309/\sqrt{\alpha}$. This estimate is in agreement with the experiments (by Floyd among others). The previous stochastic models predict a throughput $\sim 1.22/\sqrt{\alpha}$. If the constants 1.309, 1.22 are not very different, substantial differences concerning the qualitative behavior of the protocol have been shown.

- A global study of TCP [16]. A network with N nodes is assumed to use TCP to transmit packet. Flows are analyzed instead of packets. In this situation, the algorithm to allocate bandwidth is represented by the MaxMin algorithm. MaxMin is very difficult to analyze or to simulate. It can be described as the resolution of an optimization problem. We investigated the Min algorithm in the case of a random linear network. If a flow interacts with at most K flows in any of the links of its route, it received the capacity $1/K$ to transmit data. This algorithm is known as a lower bound of MaxMin. In the model considered, communications are randomly established on the real line, the global throughput is investigated. This preliminary study will be completed in the future.

Communication primitives

We have produced several new results in the area of fundamental communication primitives. Such primitives have been intensively studied before because they form major building blocks for the efficient use of parallel computers. These results are remarkable since most important questions seemed to be largely closed for a decade. The results are described going from low level primitives to more complex operations. A simple randomized permutation routing algorithm for hypercube networks [174] closes a long standing factor two gap between the best known algorithm and the trivial lower bound. A similar factor two improvement for the problem of broadcasting a long message from one processor to all others is described in [173]. A natural generalization of broadcasting is the gossiping problem where every processor has a message it wants to broadcast. Here, powerful heuristics have been developed that can be applied to any specific network as well as a new algorithm for butterfly networks that breaks a long standing barrier for very large networks. A next generalization step is if every processor has different messages (of possibly different length) for each other processor. For one natural setting, [12] gives an algorithm that beats an apparent lower bound and closes the gap to the real lower bound. We have also developed a very simple all-to-all algorithm for uniform messages that works robustly on the emerging main-stream architecture of clusters of symmetric multiprocessors. The broadcasting algorithm mentioned before [173] is also of particular interest for such machines.

Static and dynamic load–balancing

We have studied the k-partitioning problem for graphs, where the task is to divide the set of vertices of a graph equally into a given number of parts while keeping the number of crossing edges between vertices belonging to different parts as small as possible. We refer to such a number of edges as the *cut size* of the partition. The special case of partitioning the graph into 2 parts is called a *bisection*, and the minimal cut size of all balanced bisections of a graph is called its *bisection width*. (In the general case *k-section width*.)

There are several results on bounds on the bisection width of regular graphs. Results for 3- and 4-regular graphs are of special interest because these are the lowest non-trivial degrees. Some previous results for small degrees have been generalized to results for larger degrees.

It is of general theoretical interest to improve previous upper bounds on 3- and 4-regular graphs. Moreover, there are some direct applications of these results. As a motivating example, upper bounds on the bisection width of 4-regular graphs have successfully been applied to the configuration of transputer systems.

In [116] new upper bounds are derived on the bisection width of graphs which have a regular vertex degree. It is shown that the bisection width of large 3-regular graphs with $|V|$ vertices is at most $(1/6) \cdot |V|$. For the bisection width of large 4-regular graphs an upper bound of $(2/5) \cdot |V|$ is obtained.

As mentioned above, there are graph-partitioning problems in a wide range of applications. For example, when executing processes on parallel computer systems inter-processor communication is a major bottleneck. One way to address this problem is to minimize the communication between processes that are mapped to different processors. This translates to the k-partitioning problem of the corresponding process graph, where k is the number of processors. The classical spectral lower bound of $\frac{|V|}{2k} \sum_{i=1}^{k} \lambda_i$ for the k-section width of a graph is well-known. In [89] new relations between the structure and the eigenvalues of a graph are shown and a new method to get tighter lower bounds on the k-section width is presented. This method makes use of the level structure defined by the k-section. The authors define some global expansion property and prove that for graphs with the same k-section width the spectral lower bound increases with this global expansion. They also present examples of graphs for which our new bounds are tight up to a constant factor.

In [88], we consider the graph-partitioning problem by relaxing the balancing constraint, i.e. they have simply an upper bound for the number of vertices belonging to each partition. They present new and generalized lower bounds for the k-section width within this model. These bounds base on the well known lower bound of embedding a clique into the given graph with minimal congestion. This is equivalent to a multicommodity flow problem where each vertex sends a commodity of size one to every other vertex. The new bounds use arbitrary multicommodity flow instances for the bound calculation, the critical point for the lower bound is the *guaranteed cut flow* of the instances. Furthermore, a branch&bound procedure based on these bounds is presented and finally it is shown that the new bounds are also useful for lower bounds on classes of graphs, e.g. the Butterfly and Benes graph.

Another research topic within the last year is based on the relation between the properties of a graph and its suitability as interconnection network. One of the fundamental properties of a graph is the number of distinct eigenvalues of its adjacency or Laplacian matrix. Determining this number is of theoretical interest and also of practical impact. Graphs with small spectra exhibit many symmetry properties and are well suited as interconnection topologies. Especially load balancing can be done on such interconnection topologies in a small number of steps. In [86], we are interested in graphs with maximal degree $O(\log n)$, where n is the number of vertices, and with a small number of distinct eigenvalues. Their goal is to find scalable families of such graphs with polylogarithmic spectrum in the number of vertices. They also present the eigenvalues of the Butterfly graph.

Thread scheduling

In [77], an efficient, randomized, online, scheduling algorithm for a large class of programs with write-once synchronization variables is presented. The algorithm combines the work-stealing paradigm with the depth-first scheduling technique, resulting in high space efficiency and good time complexity. By automatically increasing the granularity of the work scheduled on each processor, the algorithm achieves good locality, low contention and low scheduling overhead, improving upon a previous depth-first scheduling algorithm [Blelloch et al. '97] published in SPAA'97. Moreover, it is provably efficient for the general class of multithreaded computations with write-once synchronization variables, improving upon algorithm DFDeques (published in SPAA'99 [Narlikar '99]), which is only for the more restricted class of nested parallel computations.

More specifically, consider such a computation with work T_1, depth T_∞ and σ synchronizations, and suppose that space S_1 suffices to execute the computation on a single-processor computer. Then, on a P-processor shared-memory parallel machine, the expected space complexity of our algorithm is at most $S_1 + O(PT_\infty \log(PT_\infty))$, and its expected time complexity is $O(T_1/P + \sigma \log(PT_\infty)/P + T_\infty \log(PT_\infty))$. Moreover, for any $\epsilon > 0$, the space complexity of our algorithm is $S_1 + O(P(T_\infty + \ln(1/\epsilon)) \log(P(T_\infty + \ln(P(T_\infty + \ln(1/\epsilon))/\epsilon))))$ with probability at least $1 - \epsilon$. Thus, even for values of ϵ as small as e^{-T_∞}, the space complexity of our algorithm is at most $S_1 + O(PT_\infty \log(PT_\infty))$ with probability at least $1 - e^{-T_\infty}$. These bounds include all time and space costs for both the computation and the scheduler.

Data management in networks

We have studied static and dynamic data management problems in computer systems connected by networks. A basic functionality in these systems is the interactive use of shared data objects that can be accessed from each computer in the system. Examples for these objects are files in distributed file systems, cache lines in virtual shared memory systems, or pages in the WWW. In the static scenario we are given read and write request frequencies for each computer-object pair. The goal is to calculate a placement of the objects to the memory modules, possibly with redundancy, such that a given cost function is minimized. In detail, the following three problems were considered.

A hierarchical bus network $T = (V, E)$ uses hierarchically, tree-like connected buses as a communication network. New communication technologies like SCI (Scalable Coherent Interface) make such networks very attractive, because they allow their easy construction and guarantee reasonable communication performance. Such networks can be modelled as tree networks: leaves correspond to processors, inner nodes to buses, edges to switches, and bandwidths of inner nodes and edges are related to bandwidths of buses and switches, respectively.

For these hierarchical bus networks, we investigate in [48] the static data management problem with the goal to minimize the congestion, i.e., the maximum over the load of all edges and inner nodes, induced by the read and write frequencies, divided by the bandwidth of the edge or inner node, respectively. It is known that this problem can be solved optimally in linear time if inner nodes are allowed to hold copies of shared data objects. In our model, inner nodes correspond to buses and therefore cannot store copies of shared data objects. It is shown that this restriction increases the complexity of the placement problem drastically: It becomes NP-hard. On the other hand, the main contribution of [48] is an algorithm that calculates a 7-approximation of all objects in X on a hierarchical bus network $T = (V, E)$ in time $O(|X| \cdot |V| \cdot \text{height}(T) \cdot \log(\text{degree}(T)))$.

With the widespread use of commercial networks, as, e.g., the Internet, it is more and more important to consider commercial factors within data management strategies. The goal in previous work was to utilize the available resources, especially the bandwidth, as good as possible. We present data management strategies for a model in which commercial cost instead of the communication cost is minimized, i.e., they consider a metric communication cost function and a storage cost function.

In [71], we introduce new deterministic algorithms for the static data management problem on trees and arbitrary networks. Their algorithms aim to minimize the total cost. This is the first analytic treatment of this problem which is NP-hard on arbitrary networks. The main result is a combinatorial algorithm that calculates a constant factor approximation for arbitrary networks in polynomial time. Further, the authors present an algorithm for trees that calculates an optimal placement of all objects in X on a tree $T = (V, E)$ in time $O(|X| \cdot |V| \cdot \text{diameter}(T) \cdot \log(\text{degree}(T)))$.

An important aspect of dynamic data management is data tracking. Consider an arbitrary distributed network in which large numbers of objects are continuously being created, replicated, and destroyed. A basic problem arising in such an environment is that of organizing a distributed directory service for locating object copies. In [97], we present a new data tracking scheme for locating nearby copies of objects in arbitrary distributed environments. Their tracking scheme supports efficient accesses to data objects while keeping the local memory overhead low. In particular, the proposed tracking scheme achieves an expected $polylog(n)$-approximation in the cost of any access operation, for an arbitrary network. The memory overhead incurred by this scheme is $O(polylog(n))$ times the maximum number of objects stored at any node, with high probability. It is also shown that the new tracking scheme adapts well to dynamic changes in the network.

BSP algorithms

We have focused on two main activities in the context of bridging-models for parallel computing, in particular, in the investigation of the so-called BSP (bulk-synchronous parallel machines) model. On the one hand, they develop and analyze BSP algorithms for combinatorial problems and they investigate the impact of the model used on the design of these algorithms, on the other hand, they implement a programming environment where one can actually execute BSP programs efficiently.

They developed an algorithm for the broadcast problem [70] that obeys locality constraints of the parallel machine where it is intended to be executed. Let the machine have p processors numbered from 1 through p, and assume that the communication time between two processors i and j depends only on their distance $|i - j|$. E. g., it is conceivable that on the hypercube this time is $l(2^i) = O(i)$ if the numbering of the processors forms a Hamiltonian cycle. Under this locality assumption, the authors designed an algorithm for broadcasting an item initially stored in processor 1 that works as follows. First, a broadcast is executed only on the processors 1 through $\sqrt{p}$. Then processor i that already has an item sends it to processor $i \cdot \sqrt{p}$. Finally, on all "canonical" contiguous blocks of $\sqrt{p}$ processors, a broadcast is executed. In [70] it is shown that this algorithm is faster than the trivial broadcast algorithm. E. g., on a BSP computer with a hypercubic communication mechanism, the runtime of the trivial broadcast is $O((\log p)^2)$, whereas the runtime of the method mentioned is $O(\log p \log \log p)$. Furthermore, it is proved that this method is optimal for a wide range if the BSP parameters and the locality assumption.

An overview of further impacts of the used model on the design of efficient parallel algorithms is presented in [169]. Besides the connection between locality and the broadcast problem already mentioned, it is shown how the aspects "latency" and "critical block-size" influence the development of fast algorithms for solving the multisearch problem.

A platform-independent programming environment was implemented. Virtual processors were introduced such that their execution can be stopped on the current processor and resumed at a different one. This was implemented on a Linux workstation cluster. The current state of the implementation can be inspected on the web page

`http://www.upb.de/~pub/alcom-ft/` .

Prefetching and caching in the WWW

In [41], we consider the integration of caching and prefetching techniques for the access to the World Wide Web. Caching and prefetching have often been studied as separate tools for enhancing the access to the World Wide Web. The goal of this work is to show potentialities and limitations of integrating Caching and Prefetching for improving the performances of web navigation. A new prefetching technique was proposed that uses a limited form of user cooperation to establish which documents to prefetch in the local cache at the client side. It was shown that the new prefetching technique is highly beneficial only if integrated with a suitable caching algorithm. Two caching algorithms were considered, Greedy-Dual-Size and Least Recently Used, and it was demonstrated on trace driven simulation that Greedy-Dual-Size with prefetching outperforms both LRU with prefetching and a set of popular caching algorithms. A prototype of a local proxy was developed implementing our prefetching/caching technique to be used with a web browser.

Distributed algorithms

In [32], we propose a fully-dynamic distributed algorithm for the all-pairs shortest paths problem on general networks with positive real edge weights. If Δ_σ is the number of pairs of nodes changing the distance after a single edge modification σ (insert, delete, weight-decrease, or weight-increase) then the message complexity of the proposed algorithm is $O(n\Delta\sigma)$ in the worst case, where n is the number of nodes of the network. If $\Delta_\sigma = o(n^2)$, then this is better than recomputing everything from scratch after each edge modification. Up to now only a result of Ramarao and Venkatesan was known, stating that the problem of updating shortest paths in a dynamic distributed environment is as hard as that of computing shortest paths.

WP 3: Production and Transportation Planning

Participants: Cologne, Paderborn, MPI, Utrecht
Coordinating site: Utrecht
Work Package Leader: J.A. Hoogeveen and J. van Leeuwen
Number of technical reports: 7, of which 5 also appear in other work packages

WP 3 aims to explore algorithmic issues in the area of production and transportation planning. We focus on the development of algorithms for solving or finding good solutions to problems from industry, which requires both practical, engineering work and fundamental research.

The work in the package is divided into a number of subtasks. The subtasks concern the following topics: scheduling and planning; network planning; and crew assignment and rostering.

Scheduling and planning

One scheduling problem which we have been working on originates from the pharmaceutical industry. This is a problem that can be modelled as a job shop scheduling problem with the side-constraint that waiting time between the execution of some consecutive steps in the production process is not allowed, because of the instability of intermediate products. This introductory research is described in deliverable D9.

In [125], we applied a combination of column generation and Lagrangian relaxation to another scheduling problem. This turned out to be quite successful, and it is expected that its applicability can be generalized to a wider class of scheduling problems and possibly other problems as well. In [106], we analyzed a number of scheduling problems with controllable processing times, which can be shortened at the expense of using additional resources. The quality of a solution then also depends on the amount of additional resources used; the goal is to identify the trade-off curve in polynomial time, which was possible in one of the cases. We also studied a special case of a flow shop problem in which processing times are related to demanded quantities [122].

Network planning

In [13], we studied a number of constrained network optimization problems, resulting in a generic software package, called CNOP, to solve this kind of problems. The problems under consideration consist of an easy to solve basic problem that is complicated by one or more resource constraints.

To solve such a problem, first the problem is reformulated as an integer linear programming problem with a variable x_p for each possible solution p, then the integrality constraints are relaxed, the resulting linear programming problem is dualized, and solved by a so-called hull approach, which yields both a lower and an upper bound. Eventually, the possible gap between the two bounds is closed through a solution-ranking procedure. The CNOP package implements this method and can be used to solve constrained shortest path problems and constrained minimum spanning tree problems. It also can be easily adapted to solve other constrained network optimization problems.

Crew assignment and rostering

Within this task, we studied hybrid algorithms for complex crew rostering problems. These problems are highly difficult to solve in practice as the basic assignment problem is perturbed by difficult rules and regulations reflecting legal aspects, contracts, preferences, etc. Moreover, rules and regulations may vary significantly from one company to another making it thus evident to focus on general models that can in-cooperate various rule structures.

The classical column generation model was improved by ideas from constraint programming leading to a model offering high expressiveness for rules and regulations while still providing good solutions in reasonable computation time (see [91]).

Furthermore, a constraint programming based heuristic tree search algorithm was developed, which was combined with the column generation approach. Then it was investigated how both algorithms can be coupled to overcome their inherent weaknesses by such an integration (see [90]).

Interestingly, similar methods like those developed for crew scheduling were also quite successful when being applied to a multimedia problem (see [87, 100]).

WP 4: Generic Methods

Participants: Aarhus, Barcelona, Cologne, INRIA, MPI, Paderborn, CTI, Rome, Utrecht, Warwick (all sites)
Coordinating site: Warwick
Work Package Leader: Mike Paterson
Number of technical reports: 112, of which 57 also appear in other work packages

The principal aim of our work in WP 4 is to develop new algorithms and analysis techniques to exploit future and developing technologies. In many cases, the results obtained in this area support research performed within the framework of other packages. Experimental algorithmics (WP 5) is informed by detailed analyses of algorithms, which may give precise characteristics of running times or space requirements of programs (rather than less informative $O(.)$-estimates). The control of complex networks (WP 2) cannot rely on centralized information, but has to depend instead on largely unpredictable events best represented by stochastic models. In many applications, the interpretation of massive data sets (WP 1) may involve distinguishing signal from noise. Decisions are then inherently statistical and need to be based on sound probabilistic models of what is significant (unexpected) and what is not.

Randomness models and probabilistic analysis.

One component of the research carried out within WP 4 is the precise analysis of discrete models of randomness. Such studies are essential to accompany the major shift of emphasis in algorithms design from worst-case behaviour to typical or average-case performance, and the enormous importance of randomized algorithms in present-day computing.

A characteristic feature of research within ALCOM has been strength in the area of "random combinatorics" arising from algorithmic design. The methods range from probabilistic to analytic. Thanks to ALCOM work, we are seeing the emergence of what could be termed a "theory of combinatorial processes", a largely European endeavour. Connections with the more classical theory of stochastic processes continue to develop and are now capable of attacking problems like large-scale networks or very large data sets. In these areas, ALCOM teams often occupy leading positions world-wide.

On the methodological side, the 100-page report [8] establishes a very general framework for the quantitative analysis of both finite-state and context-free models. Although such models are ubiquitous in computer science, their most fundamental quantitative aspects had never before been systematically analysed. Finite models are in many ways the largest nontrivial "decidable" class of models of computer science. They arise naturally in formal verification (model checking) but, more importantly for us, their metric properties are closely related to finite Markov chains. Hence they are central to the interpretation of large textual data, and corpuses of biological sequences [43] and even the dimensioning of communication protocols like TCP (see [33] from WP 2). Context-free models are important as they are structurally richer than finite-state models; for instance, the next generation of structural models for genomics is likely to be of this sort (due to secondary structures). The report [8] is the first to provide complete decision procedures for counting problems associated with context-free models; the report [145] discusses stochastic context-free grammars, a closely related notion, and describes corresponding experiments on RNA secondary structures.

An important area of application of the general methodology developed within the project is to large databases of sequences (this is related to WP 1). To what extent is the occurrence of a pattern expected within the model, and hence without significance? Such questions are addressed in [158, 159]. The paper [43] shows for the first time how formal models combined with analytic techniques permit us to predict what to expect regarding the number of occurrences of subsequences (i.e., symbols are not necessarily consecutive), the problem being initially motivated by intrusion detection and the filtering of large traces of concurrent events in computer systems.

Networks constitute another important area of application of these methods (see also the WP 2 section of the report). First, interconnection networks tend to have a special geometry and it is of great interest to be able to take advantage of or understand the limitations of this feature. (For instance, interconnection graphs are usually nondense.) Modelling issues are attacked in [11, 63, 66], where tradeoffs between density and robustness (measured by multiple connectivity) are obtained for the first time. In the thesis [157], deep results are developed on several universal models of "geometric graphs" (see also [136]) that appear to represent well what nature provides.

The crucial role of distributed communication networks at scales ranging from local to international poses problems of achieving efficient and coordinated communication without any centralized control. New probabilistic tools are needed for the design and analysis of reliable protocols with performance guarantees to regulate the transmission and routing of information in networks. Current research within the project is investigating routing strategies [118, 120], data management [71], load balancing [72], contention resolution (as in Ethernet) [7], and stability [10], while fluid limits inspired by physics enable us to classify precisely reservation policies [44].

Finally, the subject of probabilistic analysis has strong historic ties with the fine optimization of general purpose algorithms and data structures. Analysis-based optimization proves to have far-reaching consequences in the design of layout algorithms [136, 157] and dynamic tables (see the treatment of the randomized structure of skip-lists in [154]). Studies by ALCOM researchers and others demonstrate to practitioners the superiority of such randomized structures over classical balanced tree schemes in many situations. Also, we have attained for the first time a very precise understanding of the "HAKMEM" algorithm used for robust decision of orientation in computational geometry (see [1] for the results and the unusual level of mathematical sophistication required). New perspectives in the random generation of test sets [20] also begin to benefit from fundamental studies and a joint postdoctoral position has been arranged between MPI and INRIA for 2001-2002 on this theme. Recent progress in randomness quantification makes it possible to generate uniformly at random in quasi-linear time large complex objects of interest in simulation, in the validation of algorithms and software, as well as in performance-based tuning of computer systems.

Optimization methods and on-line scheduling.

In many applications, optimization is severely limited by the on-line nature of the data. The quality of an on-line algorithm is customarily measured by its *competitive ratio*, comparing it with an (off-line) algorithm which receives the same data, but all at once. The classic on-line bin packing problem is revisited in [50], where it is shown that *admission control*, i.e., rejecting some input items even when there is space for them, improves the competitive ratio of the First-Fit algorithm.

We studied the problem of scheduling a sequence of jobs arriving on-line on a finite set of identical machines. This is a classical problem in the theory of algorithms and has many motivating applications, ranging from efficient CPU scheduling in single- or multiple-processor computers to data delivery over multiple channels in wireless data servers. Our efforts were devoted to scheduling disciplines that minimize end-user-perceived latency. The metrics we considered are average flow time and average stretch. The *stretch* of a job is the ratio of its flow time to its processing time, and measures the quality of service provided to a job in proportion to its requirements.

We considered the framework in which jobs may be pre-empted. As to the information available to the algorithm, we first considered the clairvoyant case, in which a job (i.e., its processing requirements) is completely revealed to the on-line algorithm at its release date. With respect to this framework and considering the average stretch as the performance metric, we showed that an algorithm proposed to minimize average flow time in the clairvoyant case is also $O(1)$-competitive for this problem. The algorithm for which we proved this result does not use migration, i.e., a pre-empted job may only be resumed on the same machine it was originally assigned to. We showed that preemption is essential in this problem by proving an $\Omega(\Delta)$ lower bound on the competitive ratio of any randomized online algorithms, where Δ is the ratio between the maximum and the minimum processing time of a job. These results are described in [45].

In the non-clairvoyant case, the algorithm has no knowledge of a new job's eventual processing time when it arrives. This has been a long-open problem. In particular, it is known that no on-line deterministic algorithm can be competitive for this problem, while first, partial results were proposed for the single machine case only recently. We proved that a modification of the randomized Multi-level Feedback algorithm (RMLF), first proposed by Kalyanasundaram and Pruhs for the single machine case, is $O(\log n \log \frac{n}{m})$-competitive on m machines against the oblivious adversary. We also showed that the same RMLF algorithm achieves an $O(\log n)$-competitive ratio against the oblivious adversary on a single machine, which is also the lower bound.

Both the Windows NT and the Unix operating system scheduling policies are based on the Multi-level Feedback algorithm, thus our results also provide one theoretical validation of an idea that has been very effective in practice over the last two decades. All these results were presented in [46].

As an extension, we have started studying the most general objective function that includes average flow time and average stretch as a special case, i.e., weighted flow time. In this case, each job comes with a positive weight. The weighted flow of a job is defined as the flow time of the job times its weight. The goal is minimizing the total weighted flow of the jobs. Little is known about this problem and only very recently were preliminary results proposed. We studied a special case of this problem, which we called *deadline scheduling*. Here, all jobs are released at time 0 and the goal is minimizing the overall weight of the jobs that have not been completed by an unknown deadline t. We have a constant competitive algorithm for the problem, and showed that an on-line algorithm is constant competitive with respect to minimizing weighted flow time if and only if it is constant competitive with respect to deadline scheduling for any t. These results are given in [42].

Another complication is considered by [30]: there are *two different* cost measures on jobs. The performance of an algorithm is compared to an off-line adversary which is optimizing with respect to its own combination of the two criteria.

When combined, Operations Research and Constraint Programming methods can produce a powerful optimization algorithm. The effect of integrating shortest path algorithms and complex rule systems for Crew Assignment Problems are investigated in [90, 91]. Where one branch of investigation considered an improved classical column generation approach only, the other branch combined two distinct approaches, with the inherent weaknesses of each approach being removed by the advantages of the other (see WP 3).

For a multimedia application – automatic tape recording in a digital TV context – a new stable set algorithm for co-interval graphs was combined with some fast variable-fixing algorithms for knapsack problems [100]. Both algorithms use a preprocessing phase that consists of sorting some data in time $O(n \log n)$. Afterwards, each node of a tree search only needs time $O(n)$ for fixing variables. Thus the average processing time is $O(n)$ if more than $\Omega(\log n)$ nodes are visited within the tree-search. Compared to previously known algorithms this idea decreases the running time from $O(n \log n)$ per node to $O(n)$ for knapsack problems, while preserving the solution quality.

Approximation methods.

One of the most popular methods in approximate counting algorithms is the Markov Chain Monte Carlo method, which exploits the fact that the problem of approximately counting combinatorial structures can often be reduced to the problem of sampling the structures almost uniformly at random. Samples are produced by simulating an appropriate Markov chain, and the chief difficulty is bounding the *mixing time* of the chain (the amount of simulation required before the stationary distribution of the chain is (nearly) reached). General sampling (and counting) algorithms for sampling orbits are studied in [4]. For example, these algorithms can be used to approximately count unlabelled combinatorial structures. A specific application of the MCMC method is studied in [5], in which a Markov chain modelling the iterated Prisoner's Dilemma game is shown to be rapidly mixing.

A classical problem in computer science is to verify if a given object possesses a certain property or not. For example, we want to decide whether a Boolean formula is satisfiable, or whether a graph is connected. The goal of *property testing* is to provide answers to a relaxed form of this classical decision problem. Instead of solving the problem exactly, we just want to distinguish between the case when the object has a certain property and the case when the object is 'far away' from that property. Thus we can formulate the goal of property testing as follows:

> Let O be an unknown object from a class of objects C, and Q a fixed property of objects from C. Determine (possibly probabilistically) if O has property Q or if it is far from any object in C which has property Q, where distance is measured with respect to some distribution D on C.

Report [133] analyses a testing algorithm for ℓ-colouring of k-uniform hypergraphs. While the general decision problem is known to be $\mathcal{NP}$-hard, a testing algorithm is presented that can distinguish between the case when the hypergraph has a proper ℓ-colouring and the case when it is ϵ-far from having such a colouring, in time independent of the size of the hypergraph (a hypergraph is ϵ-far from having a proper l-colouring, if we have to change an ϵ-fraction of the entries in its k-dimensional adjacency matrix for it to have a proper ℓ-colouring).

In [132], property testing is used to maintain approximate data structures for sets of moving points with unpredictable motion. The property testing algorithm checks whether the data structure is 'far from correct'. If it is, the algorithm finds an error in the structure and repairs it. Such data structures are called *soft kinetic data structures*. A competitive analysis of the quality of a soft kinetic data structure is given with respect to the dynamics of the system. Soft kinetic data structures are presented for maintaining classical data structures: sorted arrays, balanced search trees, heaps, range trees and Euclidean minimum spanning trees.

Graph and network algorithms

Part of our effort was devoted to studying dynamic path problems on directed graphs. We have introduced a general framework for casting fully dynamic transitive closure into the problem of reevaluating polynomials over matrices. Using this technique, we have devised a new deterministic algorithm that improves the best known bounds for both fully-dynamic and deletion-only transitive closure on general directed graphs [147, 25]. Our algorithm maintains explicitly the transitive closure as a Boolean matrix, supporting updates in $O(n^2)$ amortized time and queries in unit worst-case time. Since an update may change as many as $\Omega(n^2)$ entries of this matrix, this seems to be the best update bound that one could hope for over this class of algorithms. In the case of deletions only, we can perform updates faster in $O(n)$ amortized time. Using the same matrix-based approach, we have also considered the problem of maintaining implicitly the transitive closure of a directed graph, devising the first algorithm that performs both updates and reachability queries in subquadratic time per operation [147, 25]. This result proves that it is actually possible to break through the $O(n^2)$ barrier on the single-operation complexity of fully dynamic transitive closure. Our subquadratic algorithm is randomized Monte Carlo and supports updates in $O(n^{1.58})$ and queries in $O(n^{0.58})$ worst-case time. Another problem that we have considered is how to maintain shortest paths in dynamic graphs. As a main contribution, we have presented the first fully dynamic algorithm for maintaining all pairs shortest paths in directed graphs with real-valued edge weights [150]. In particular, we have shown how to support updates in a graph whose edges can assume at most S different real values in $O(S \cdot n^{2.5} \log^3 n)$ amortized time. Our algorithm is deterministic and can answer queries about shortest paths in optimal worst-case time. Although studied since the late 60's, no previous fully dynamic algorithm was known for this problem. In the special case where edge weights can only be increased, we have also devised a randomized Monte Carlo algorithm that supports updates faster, in $O(S \cdot n \log^3 n)$ amortized time.

In many applications, efficient algorithms can take advantage of the restricted nature of the graphs and networks which arise, as parameterised by notions related to treewidth. One such notion is that of cutwidth, for which linear time algorithms have been established when the cutwidth is small [114], and polynomial time algorithms when the treewidth and degree are small [115]. The treewidth is used to obtain new results for computing the minimum dominating set and related problems on planar graphs [112]. Experiments have been carried out to compute the treewidth of networks, arising from applications from decision support systems (probabilistic or Bayesian networks) and telecommunication, where it is shown that preprocessing with graph reduction is very useful towards computing the treewidth of these networks (as is required for these applications). Also, efficient algorithms have been found for the problem of finding for a given graph a regular supergraph of specified degree with as few additional vertices as possible [122].

Distributed computing.

We have considered two approaches to the problem of transforming sequential programs into BSP and distributed versions.

In [156], a collection of algebraic laws describing the behaviour of BSP programs is developed. Asynchronous communications and *barriers* are treated as a generalised form of multiple assignment. The derivation, from its sequential representation, of a parallel BSP-style version of Floyd's shortest path algorithm is used to illustrate the use of the proposed algebraic laws.

In [139], an experiment is conducted to assess the feasibility of transforming BSP computations into systems employing explicit asynchronous *two-sided* communications. The aim is to show how a verified BSP program may be transformed into a guaranteed correct conventional message-passing implementation.

Complexity bounds.

Trade-offs between one resource bound and another provide insight into the intrinsic complexity of computational problems. Results of this nature obtained in the first year of ALCOM-FT illustrate this well.

For comparison-based sorting of n w-bit integers, the product of time and space required is known to be $\Theta(n^2)$, and Beame has shown that the same lower bound holds for non-comparison-based sorting too. In the latter case, no matching upper bound was known for time below $\Theta(n \log n)$. Results in [141] match the lower bound for any time down to $\Theta(n(\log\log n)^2)$, or even down to $\Theta(n \log\log n)$ with high probability using a randomized algorithm.

The R-way branching programs are a restricted computational model. In [172], the problem of deciding whether any pair out of n elements are "close" to each other is shown to have a time-space trade-off, extending a recent result by Ajtai (STOC'99).

In external memory (or I/O) models, it is the number of I/O operations, accesses to external memory, which replace computation time. A general technique for deriving I/O-space trade-offs from corresponding time-space trade-offs is developed in [71] and demonstrated with I/O-space trade-offs for sorting and the "element distinctness" problem.

The complexity of *counting* dates back to the 1970's with the introduction of Valiant's class #P. While much is known about the complexity of exact counting, *approximate* counting is much less understood. We do know that approximate counting is significantly easier than its exact counterpart (approximate counting is not much harder than NP, whereas #P-hard problems are as hard as the whole polynomial hierarchy). The paper [6] studies the complexity of approximate counting using the vehicle of approximation-preserving reduction. It describes and investigates three classes of counting problems that are interreducible under approximation-preserving reductions: the class of problems that admit an FPRAS, the hardest problems in a logically-defined subclass of #P, and the class of problems which are complete for #P with respect to approximation-preserving reductions.

An active research activity in recent years has been to tighten upper and lower bounds on the clause-to-variable ratio for random 3-SAT formulae which marks the abrupt change from almost all being satisfiable to almost all being unsatisfiable. The latest result in this competition is the upper bound of 4.571 in [103].

New data structures.

Dynamic tables that support search, insert and delete operations are fundamental and well-studied in computer science. There are many well known data structures that can be used for such *dictionaries*, including binary search tree (BSTs, for short), skip lists and tries among others. Since BSTs perform poorly when they are skewed, many variants of balanced BSTs have been devised. The method presented in [167] achieves logarithmic worst-case cost by using the sizes of the subtrees as balancing information. Therefore, unlike AVL trees and red-black trees, rank operations are efficiently supported without any changes in the data structure. Compared to weighted binary search trees, the operations involved with our method are likely to be less costly in most real situations. The paper presents empirical evidence that our trees are faster than weighted BSTS, and as fast as other variants of BSTs.

Relaxed balance is the term used for search trees where the rebalancing has been uncoupled from the updating, such that they become processes which can be controlled separately. One advantage is that rebalancing can be turned off in order to speed up the processing when update requests come in bursts. When the burst is over, the tree can be rebalanced again, while searching and updating still continue. Another motivation comes from using search trees on shared-memory architectures. There, locking of nodes is necessary when rebalancing the tree, but the locking limits the amount of parallelism possible. In search trees with relaxed balance, rebalancing operations can be carried out such that the locks are very localized in time and space—in particular, exclusive locking of whole paths or step-wise exclusive locking down paths can be avoided. In [128], a search tree with relaxed balance having a height bound of $c \cdot \log n$ for any $c > 1$ is presented. The previous best bound was $c = 1.44$. In [51], the repertoire of B-trees with relaxed balance is extended to include group updates, and in [52], variants of B-trees with relaxed balance having improved space utilization is presented.

The classic data structures for dictionaries use $O(\log n)$ time for searches, inserts and deletes, matching the information theoretic lower bound when access probabilities are uniform i.i.d. If access probabilities are non-uniform but still i.i.d., there are other weighted data structures such as D-trees, biased search trees, splay trees and treaps which can achieve optimality. In many applications, however, the source of non-uniformity in access probabilities is locality of reference: examples include memory, cache, disk and buffer management and emerging applications in internetwork traffic management. In such applications, the access probability of any given key is not i.i.d., but decreases with idle time since the last access to the key. In [153, 154], a *lazy update* scheme is introduced to deal with such distributions, and is illustrated by a new data structure, the biased skip list (BSL). More specifically, let $r(k)$ be the number of distinct keys accessed since the last access to key k i.e., $r(k)$ is the move-to-front rank of k. Let $r_{\max}(k)$ be the maximum rank of k during its lifetime. The lazy update scheme provides search in $O(\log r(k))$ time and insert/delete in $O(\log r_{\max}(k))$ time. These bounds are optimal. The practicality of BSL is demonstrated with experiments on real and synthetic data with various degrees of bias.

Hashing covers another important set of techniques for the dictionary problem. We have contributed to both theoretical and practical aspects of hashing. A new deterministic algorithm [68] could be the first step towards hashing schemes that are in all respects better than traditional (comparison-based) techniques for simple search. Another paper [67] provides a precise characterization of the optimal behaviour of hashing schemes with respect to the number of memory locations accessed in the process, an important performance parameter in modern computing systems. Cuckoo Hashing [124], a practically promising hashing algorithm, was recently developed, and is now under consideration for inclusion in the LEDA software library. Finally, a relaxation of a standard search problem has been considered, trading slightly imprecise storage for an unseen combination of space and time efficiency [126].

SCIL (Symbolic Constraints for Integer Linear Programming)

SCIL is a software project being developed jointly by the MPI and Cologne sites and by Alexander Bockmayr (Nancy) and Thomas Kasper (SAP). Our aim is to design a high-level description language for combinatorial optimization problems, and to deliver an easy-to-use software system for such problems. The algorithms are to be tested on recognised test instances and problems supplied by our industrial contacts.

During the first year of ALCOM-FT, the design of the high-level description language went through several iterations and resulted in a fifty-page internal document. The implementation in C++ similarly went through several iterations. It now consists of a realization of the high-level description language, an interface to ABACUS and LEDA, and a small library of constraints and their separation and pricing routines. The implementation is prototypical and, at present, can only be used by the developers. However, the system has been successfully used to develop optimization algorithms for problems in curve reconstruction and computational biology.

WP 5: Experimental Algorithmics

Participants: Aarhus, Barcelona, Cologne, INRIA, MPI, Paderborn, CTI, Rome
Coordinating site: MPI
Work Package Leader: Peter Sanders
Number of technical reports: 33, of which 21 also appear in other work packages

Algorithm engineering has become an important part of the research of several ALCOM-FT sites. MPI organized the *Workshop on Algorithm Engineering* (WAE) in 2000, Aarhus organizes it in 2001. WAE is the main European event on Algorithm Engineering. Giuseppe F. Italiano, Rome, was Guest Editor of a Special Issue devoted to "Algorithm Engineering" in the journal *Algorithmica* (vol. 30, no. 4, August 2001). This special issue covers a broad range of current research in the area. He was also invited to give a plenary lecture on Experimental Algorithmics [24] at the *25th International Symposium on Mathematical Foundations of Computer Science* (MFCS 2000), held in Bratislava, Slovakia, August 28 – September 1, 2000. An invited talk on algorithm engineering was also give by Kurt Mehlhorn on the HERCMA Conference in Athens. There is also an upcoming Summer School on Algorithm Engineering organized by the Rome site, to be held on September 10–12, 2001 (`www.info.uniroma2.it/~italiano/School`).

Below we report in more detail on a Dagstuhl seminar on experimental algorithmics organized co-organized by E. Meineche Schmidt from Aarhus.

The 33 ALCOM-FT reports [1, 18, 20, 22, 23, 24, 26, 27, 28, 40, 52, 56, 58, 61, 62, 73, 74, 80, 81, 83, 99, 101, 102, 119, 124, 134, 136, 147, 148, 149, 154, 157, 165] relating to WP 5 have been authored or co-authored by members of the sites. Until now, 23 of these have appeared in conference proceedings or journals, or are scheduled to do so soon. About half of these publications are for workshops and journals dedicated to algorithm engineering (WAE,ALENEX, ACM JEA). One publication [134] was on SIGGRAPH, a highly competitive non-theoretical graphics conference.

Methodology

Experimental algorithmics is at the same time a young and an old discipline. On the one hand, there have always been people who evaluated algorithms experimentally with more or less success. On the other hand, there is recently an increasing awareness among theorists, that theoretical analysis has to be complemented with experiments more often. To do this in a meaningful way we need to answer question like

- What are relevant experiments?
- What can be learned from experiments?
- What is a good benchmark test set?
- What is a good experimental paper?

Before trying to invent new answers to these questions, one should first collect known answers. With that aim, we participated in organizing a workshop on *Experimental Algorithmics* held at Schloß Dagstuhl, Wadern, Germany. The workshop took place in the period September 10 to September 15, 2000, and was initiated and organized by Rudolf Fleischer of ALCOM-FT site MPI, and Erik M. Schmidt of ALCOM-FT site Aarhus, in cooperation with Bernard Moret of University of New Mexico, Albuquerque. In all, 45 researchers with affiliations in Austria, Canada, Denmark, France, Germany, Great Britain, Greece, Hong Kong, Italy, the Netherlands, Spain, and the USA participated in the meeting. Three invited keynote speakers, Jon Bentley, David S. Johnson, and Kurt Mehlhorn, gave one-hour position talks. The remaining 26 presentations given by participants of the meeting covered a wide range of topics in experimental algorithmics. The report of the seminar contains the abstracts of most of these presentations, as well as a summary of the plenum discussions held, and a list of participants.

To help disseminate the knowledge gained during the workshop, papers on experimental algorithmics were solicited from the participants. These papers will form a set of proceedings from the workshop, for which world-wide publication as a volume in the series *Lecture Notes in Computer Science* from Springer-Verlag, Germany has been arranged. Almost half of the articles in the volume will be co-authored by researchers from consortium members. Some of them are already available in the report database [24, 80, 81, 149]. For example, in [80] we discuss in detail how asymptotic analysis and experiments can coexist.

Experimental Evaluation of Data Structures

We use the very fundamental domain of data structures to exemplify why experiments can complement the mathematical analysis of algorithms.

A simple, robust, efficient and well known way to maintain a dictionary are (a, b)-trees. A new variant that relaxes balancing conditions [52] achieves an improved space utilization. Faster dictionaries are possible if the accesses are distributed in a nonuniform way. Biased Skip Lists [153] allow the maintenance of dictionaries so that access to frequently accessed items is very fast without incurring a big overhead for less frequently accessed items. Experiments are important in [52, 153] and many other papers to demonstrate practicability of algorithmic measures. If we can use hash functions on the elements, we can even solve the dictionary problem with constant worst case lookup and deletion time and constant expected insertion time using the new, very simple Cuckoo Hashing scheme [124]. Here we get a new reason for using experiments. Although a theoretical analysis is available for powerful hash functions, the practically important case of very simple hash functions is an open theoretical problem. This gap can be partially closed by demonstrating experimentally that simple hash functions also work well. Remarkably similar techniques (exploiting duplicate random choices in a globally optimal way) can be used for scheduling parallel disks [99]. Here, no analytical results are known. Even much more rudimentary scheduling strategies have so far escaped an analysis for the case of continuous streams of disk accesses.

Experiments are also important because we have to use very simple machine models for mathematical algorithm analysis. Hence experiments are important for validating simplifying assumptions in the model. A typical example is the priority queue data structure based on Sequence Heaps [83]. The new algorithm is first analyzed using the assumption that we have full control over the cache content. Then an implementation indicates that it is the fastest known comparison based priority queue. To arrive at such a claim one has to be very careful with the implementation. In particular, too often a new algorithms proposed by the author is faster just because it is implemented more carefully than a previous algorithm. The best way to avoid this is to implement all compared algorithms very carefully. In [83] inspection of the generated assembler code is used to argue that all implementations execute a close to minimal number of instructions. Measurements on four different processor families add further robustness to the claims.

Experimental Graph Algorithmics

The weighted matching algorithms in LEDA were re-implemented and fine-tuned . An experimental comparison of our new implementation of the general weighted matching algorithm (contained in LEDA-4.2 or higher) with the most efficient code currently available showed that our matching algorithm is considerably faster on almost all instances. A new heuristic as described in [18] was implemented for the bipartite weighted matching algorithm. Experiments were performed on different families of random instances. The experimental findings indicate a running time improvement of up to a factor of 10. The improved code is part of LEDA-4.3.

New coloring techniques based on local search have been developed. These have been implemented in C, tuned, and tested experimentally. The practical performance of our techniques on commonly used benchmarks for coloring has been very encouraging [28]. Variations of those techniques have been successfully deployed for University Examination Timetabling [27]: our algorithmic codes will be further tuned so as to be included in a timetabling server at an existing university.

We have also addressed the one-sided crossing minimization problem (CP), studying its strong relationship with the problem of computing minimum feedback arc sets in directed graphs [148]. In particular, we have devised a new approximation algorithm for minimizing crossings (PM) based on this relationship. We have experimentally compared PM with previous well-known algorithms and with recent attractive strategies. Experiments were carried out on different families of randomly generated graphs, on pathological instances for CP, and on real test sets. Performance indicators included both number of edge crossings and running time, as well as structural measures of the problem instances. The experiments clearly separated the behavior of the algorithms, proving the effectiveness of PM and showing tradeoffs between quality of the solutions and running time.

Experiments on Dynamic Graph Algorithms

We put significant effort in the experimental study of several dynamic graph algorithms. In particular, in [102] we have implemented and engineered dynamic transitive closure algorithms by Italiano, Yellin, Cicerone *et al.*, and two recent randomized algorithms by Henzinger and King. We have proposed a fine-tuned version of Italiano's algorithms as well as a new variant of them, both of which were always faster than any of the other implementations of the dynamic algorithms. We have also considered simple-minded algorithms that were easy to implement and likely to be fast in practice. We have tested and compared the above implementations on random inputs, on non-random inputs that are worst-case inputs for the dynamic algorithms, and on an input motivated by a real-world graph.

Another dynamic problem that we have investigated experimentally is fully dynamic single-source shortest paths in digraphs with arbitrary (negative and non-negative) arc weights [23]. We have implemented and tested several variants of the theoretically fastest fully dynamic algorithms proposed in the literature, plus a new variant devised to be as simple as possible while matching the best worst-case bounds for the problem. All the considered dynamic algorithms were faster by several orders of magnitude than recomputing from scratch with the best static algorithm. The experiments also revealed that, although the simple variant that we suggested is likely to be the fastest in practice, other dynamic algorithms proposed in the literature might be preferable in some cases.

Web-Based Testbed for Experimental Graph Algorithmics

A web-site functioning as a testbed for various graph algorithms has been set up (see the address `www.mpi-sb.mpg.de/~schaefer/MLLB/index.html`). The web-site can be used to submit problem instances as well as problem generators for the following graph algorithms: all variants of matching algorithms (i.e. bipartite cardinality matching, bipartite weighted matching, general cardinality matching and general weighted matching) and a max-flow algorithm. Further graph algorithms, such as the min-cost flow algorithm, will be added in the near future.

The major motivation behind setting-up this web-site is to establish a set of both benchmark instances and problem generators for various graph algorithms. Moreover, the need for making graph algorithms more accessible to the public is achieved by an easy-to-use web-interface.

Submitted problem instances are run on a compute server of the MPI and the result is reported back to the sender via e-mail. The web-site thereby keeps track of the encountered worst-case instances for each algorithm. One may download these instances from the web-site.

Additionally, one may submit worst-case generators (as C++ source files). A submitted generator is accepted only if it follows a certain specification, which is given on the web-site, and if it passes a functionality test (which for the time being and probably also in the future cannot be done automatically). Each accepted problem generator is added to the web-site and can be downloaded.

Algorithm Visualization

We have realized an algorithm animation system, called Catai (for Concurrent Algorithms and data Types Animation over the Internet) [40]. Catai is heavily based on an Object-Oriented animation approach that allows a programmer with an algorithmic background to transparently and efficiently animate a target algorithm. Multiple users can simultaneously watch and interact with a single Catai-animation using a light-weight Java animation client. We believe this to be a good compromise between two different viewpoints: the programmer's perspective, which typically includes the goal of animating efficiently and unobtrusively a given algorithmic code, and the user's perspective, which can benefit from interactive, easy-to-use, distributed and cooperative interfaces.

Another software contribution is Leonardo, an integrated environment for developing, executing, and visualizing C programs [26]. Leonardo features two major improvements over a traditional IDE. First, it provides a mechanism for visualizing computations graphically as they happen by attaching in a declarative style graphical representations to key variables in a program. As a second main feature, Leonardo includes the first run-time environment that supports fully reversible execution of C programs. The system is distributed with a collection of animations of more than 60 algorithms and data structures including approximation, combinatorial optimization, computational geometry, on-line and dynamic algorithms. Leonardo has been widely distributed on CD-ROM in computer magazines and is available for download in several software archives over the Web. It has received several technical reviews and more than 18,000 downloads over the last two years.

Part of our effort was also devoted to understanding in more details the role of visualization in algorithm engineering [149]. In particular, we have considered the main approaches, existing tools and relevant situations where visualization systems can help developers gain insight about algorithms, test implementation weaknesses, and tune suitable heuristics for improving the practical performance of algorithmic codes.

Interactive Rendering of Highly Complex Scenes

A walk-through system is a visualization of a three dimensional geometric scene where a visitor of a scene can walk to arbitrary positions. The computer has to render a new image for every new position. In order to get a smooth visualization and a good navigation, at least 20 images per second are necessary. The time for an image computation mainly depends on the number of triangles and the number of computed pixels. Therefore, for very complex scenes a walk-through in real-time is not possible. Rendering algorithms are necessary for a reduction of the scene complexity.

In [134] we present a new output-sensitive randomized rendering algorithm, the randomized z-buffer algorithm. It renders an image of an arbitrary three-dimensional scene consisting of triangular primitives by reconstruction from a dynamically chosen set of random surface sample points. This approach is independent of mesh connectivity and topology. The resulting rendering time grows only logarithmically with the numbers of triangles in the scene. We were able to render walk-throughs of scenes of up to 10^{14} triangles at interactive frame rates. Automatic identification of low detail scene components ensures that the rendering speed of the randomized z-buffer cannot drop below that of conventional z-buffer rendering. Experimental and analytical evidence is given that the image quality is comparable to that of common approaches like z-buffer rendering. The precomputed data structures employed by the randomized z-buffer allow for interactive dynamic updates of the scene. Their memory requirements grow only linearly with the number of triangles and allow for a scene graph based instantiation scheme to further reduce memory consumption.

Engineering Distributed Algorithms

Experiments for distributed algorithms have some special characteristics: they need either a scalable real network environment or a platform supporting a simulated distributed environment. Also, arbitrary asynchrony is hard, if not impossible, to implement and hence the best way to proceed is perhaps through time delay random distributions. The generation of "worst-case" instances is also an important challenge.

We are currently developing, as part of WP 2, a new platform, called Distributed Algorithms Platform (DAP), implemented in C++ using LEDA. The major goal of DAP is to provide a homogeneous environment for the simulation of distributed algorithms, regardless whether they are designed for fixed networks or for mobile networks.

We have used early versions of this platform to perform experiments for communication in ad-hoc mobile networks. In particular, [61] proposes, theoretically analyses and experimentally validates a new and efficient protocol for communication in ad-hoc mobile networks. The protocol exploits the co-ordinated motion of a small part of the network (the "support") in order to provide to various senders and receivers an efficient support for message passing. The protocol has been tested for three classes of graphs (grids, random graphs and bipartite multi-stage graphs), each abstracting a different "motion topology". These ideas have been extended in [62] where a new model of ad-hoc mobile networks, which is called hierarchical, has been introduced. Such networks are comprised of dense subnetworks of mobile users interconnected across access ports by a very fast backbone. In [119], the basic "support" idea in [61] where the support has a "snake" shape has been comparatively studied with an alternative implementation of it, that of the "runners" protocol in which the members of the support perform concurrent and continuous random walks and exchange any information given to them by senders when they meet. The experimental evaluation showed that for both protocols only a small support is required for efficient communication, and that the runners protocol outperforms the snake protocol in almost all types of inputs we considered.

In [101] a rather general model for modeling attacks in computer networks and the development of protocols for propagation of attacks under this model have been studied.

In [56, 58], the problem of routing packets in two-dimensional torus-connected processor arrays has been considered. Motivated by recent theoretical work on dynamic routing, the dynamic case where packets are continuously injected into the network is addressed. Greedy hot-potato routing algorithms are described and simulation experiments on tori of several sizes using four well-known greedy (also known as work-conserving) protocols (namely FIFO, LIFO, FTG, and NTG) for the implementation of injection buffers have been presented. The results demonstrate that there exists a greedy hot-potato routing algorithm that is stable for all greedy injection queueing protocols for injection rates very close to 100% of the network capacity. Furthermore, according to the algorithms studied, we can claim that the one-pass property is not appropriate for the dynamic case, since the system cannot achieve stability at high injection rates.

WP 6: Project Management, Dissemination, Evaluation

Participants: Aarhus, Barcelona, Cologne, INRIA, MPI, Paderborn, CTI, Rome, Utrecht, Warwick (all sites)
Coordinating site: Aarhus
Work Package Leader: Erik M. Schmidt

Project Management

The *Consortium Board*, consisting of the scientific leaders of each of the ten participating sites and chaired by the *Coordinator*, has carried the overall responsibility for the management of the project. Based on the guidelines laid down by the board, and in frequent email contact with its members, the Coordinator has performed the day-to-day actions necessary for the smooth operation of the project at large. The Coordinator has also been the point of contact between the EC and the project, and has been responsible for the submission of project deliverables and reports. A *Steering Committee* of consisting of four site leaders has been elected with the purpose of aiding the Coordinator in situations where time does not permit consulting the entire Consortium Board.

The work at each site has been organized by the site leader in question, with the help of one or more assistants. To ensure proper coordination of the activities within each work package, a *Work Package Leader* has been elected for each of the six work packages. As witnessed by the timely finalization of our deliverables, the management of the work performed has been very successful. The Work Package Leaders have also been responsible for the reporting on the work done in each work package.

The first ALCOM-FT meeting of the Consortium Board was held September 9, 2000 in Saarbrücken. Steering Committee and Work Package Leaders were elected, and details of the work to be performed were discussed. Among other things, precise responsibility for each deliverable was settled, the time, place and format of the annual review meetings were decided, and place and contents of the proposed summer schools were finalized.

In general, communication within the project, as well as between the project and the EC, has worked very well, and no problems have been encountered during the management of the project in its first year.

Dissemination

The principal means of dissemination for the ALCOM-FT project is scientific reports describing the research done. A total of 174 reports have been produced during the first year of the project. They have all appeared in the ALCOM-FT Technical Report Series which is available online.

During the first year of the project, two Summer Schools were organized. The first was held at MPI in Saarbrücken with the title *Advanced Course on the Foundations of Computer Science*, and the second was held in Paderborn with the title *ALCOM Spring School on Dynamic Algorithms*.

An ALCOM-FT website (`http://www.brics.dk/ALCOM-FT`) has been set up, containing a description of the objectives of the project, the sites participating, and the deliverables produced. The website also contains the online version of the ALCOM-FT Technical Report Series. Furthermore, several of the deliverables have separate websites describing the work performed. These can also be reached through the ALCOM-FT website. Finally, at website facilitating the sharing of lecture notes among lecturers within computer science worldwide has been initiated.

During the first year of the project, the conferences ESA, WAE, and APPROX was held at MPI, Saarbrücken, organized by the local ALCOM-FT group. All three conferences are recognized worldwide, and have subjects highly relevant for the work performed within the ALCOM-FT project. The ESA and WAE conferences were started eight and five years ago, respectively, by members of the ALCOM community, and are now among the major algorithmic European events. Additionally, many smaller conferences and workshops with subjects relevant for the project were arranged by various ALCOM-FT sites, including a Dagstuhl meeting on algorithmic engineering – see the site reports later in this document for further details.

No "competitions" involving people external to ALCOM-FT have been set up during the first project year. This will be a theme for the second year.

Evaluation

Overall, the ALCOM-FT project is in good shape. The work is progressing as planned and the productivity is high (as an example, the contract promised some 40 papers in WP 4, whereas the actual number is 112).

The international standing of ALCOM-FT is illustrated by the fact that more than 114 out of the total of 174 reports have already been published (or been accepted for publication) in the scientific community via conferences and journals, including the following key conferences: ALENEX (3 publications), APPROX (3), ESA (8), FOCS (3), ICALP (7), MFCS (4), SPAA (10), SODA (5), STACS (3), STOC (5), SWAT (5), WADS (2) and WAE (5).

The consortium maintains a high profile in PhD education which is illustrated by the fact that a total of 21 PhD's obtained their PhD-degrees during the first project year.

So far, no external review of key deliverables has taken place, mainly because most of the relevant deliverables (except from regular papers that are refereed in the usual way) were delivered at the very end of the first year.

Part II Project Partners

Aarhus

The activity at site Aarhus is based on the algorithms groups at the University of Aarhus and the University of Odense. The site is the coordinating site of the ALCOM-FT project, and is involved in Work Packages 1, 4, 5, and 6. The group in Aarhus is part of BRICS (Basic Research in Computer Science, Centre and International Research School of the Danish National Research Foundation). In total, the two groups currently comprise 12 associate/assistant professors and 10 Ph.D. students. Rolf Fagerberg and Gerth S. Brodal are paid part-time by ALCOM-FT funds; Rolf Fagerberg is assisting with the coordination of the project.

During the first year of ALCOM-FT, Jakob Pagter completed his Ph.D. with the thesis *Time-Space Trade-Offs*. Visitors to the site include Lars Arge, Sanjeev Arora, Devdatt Dubhashi, Uri Feige, Bernd Gärtner, Alessandro Panconesi, Joel Spencer, Aravind Srinivasan, Zsolt Tusa, Uri Zwick, and Anna Östlin.

Members of the site have participated in most of the important relevant conferences, including ESA, FOCS, ICALP, SODA, STOC, and SWAT. The site hosted *Workshop on Probabilistic Methods in Combinatorial Optimization* in August 2000.

The 24 ALCOM-FT reports [34, 35, 36, 37, 38, 49, 50, 51, 52, 53, 67, 68, 78, 124, 126, 128, 130, 131, 141, 144, 145, 164, 171, 172] have been authored or co-authored by members of the site. Until now, 16 of these have appeared in conference proceedings or journals, or are scheduled to do so soon.

Barcelona

The activity at site Barcelona is based on the theory group at the Department of LSI of the UPC. This first period, the site is involved in Work Packages 1, 2 and 4. In total, the group currently comprises 9 associate/assistant/full professors and 6 Ph.D. students. G. Casas and M.J. Blesa are paid part-time by ALCOM-FT funds.

During the first year of ALCOM-FT, Jordi Petit completed his Ph.D. with the thesis *Layout Problems* (see ALCOMFT-TR-01-157). ALCOM-FT visitors to the site include , Philippe Flajolet (INRIA), Paul Spirakis (CTI), Burkhard Monien (Paderborn) and Peter Sanders (MPI). Gemma Casas attended the *Workshop on Dynamic Algorithms* at Paderborn and the *Workshop on Reinforcement Learning* at London. She also visited the MPI.

The 21 ALCOM-FT reports [156, 143, 142, 139, 135, 170, 167, 166, 165, 157, 155, 152, 151, 140, 138, 136, 115, 114, 92, 39, 21] have been authored or co-authored by members of the site. Until now, 17 of these have appeared in conference proceedings or journals, or are scheduled to do so soon.

Cologne

The activity at site Cologne is based on the Institute of Computer Science at the University of Cologne. The site is involved in Work Packages 3, 4, and 5. The group currently comprises 5 Ph.D. students. Matthias Elf is paid full-time by ALCOM-FT funds.

During the first year of ALCOM-FT, Michael Jünger was a member of the program committees for COCOON2000, WAE2000, and ESA2000. Members of the site attended ISMP2000, OR2000, and GraphDrawing2000. Michael Jünger gave an invited lecture at the *8ème rencontre en programmation mathématique* in Han-sur-Lesse, Belgium. He also gave an industrial talk in the Chamber of Industry and Commerce of Cologne. Matthias Elf and Michael Jünger visited MPI.

The site organized two workshops, the *Dagstuhl Seminar on Algorithmic Techniques in Physics* in February 2001, and the *Fifth Combinatorial Optimization Workshop Aussois 2001 "Eureka, you shrink!"* in the honor of Jack Edmonds in March 2001.

The ALCOM-FT reports [84, 123] have been authored or co-authored by members of the site.

INRIA

The activity at site INRIA is based on the Algorithms Group at INRIA Rocquencourt, France. The group comprises 5 permanent members (one junior and 4 senior members), 5 PhD students, and 4 associate members (from University of Caen and Paris). The research of 2 of the PhD students (M. Durand and C. Banderier) takes place entirely within the framework of the ALCOM Project. The work packages concerned are WP 1, WP 2, WP 4, and (to a smaller extent) WP 5.

Research produced in the group continues to be largely directed towards general methods aiming at precisely quantifying randomness in large discrete structures. The framework developed is that of analytic combinatorics and it is finding applications in data structures and algorithms (WP 4 and WP 5), massive data sets (WP 1) and large interconnection networks (WP 2). We have developed an extensive visitor program (with publication of seminar proceedings [110], 150 pages) and are hosting in 2001 the *Seventh Seminar on Analysis of Algorithms*.

During the first year of ALCOM-FT, Cyril Banderier completed his Ph.D. thesis *Analytic Combinatorics of Paths and Maps* (in French). Marianne Durand joined the group as a starting PhD student working on combinatorial models. Joint research has been developed with Patras (interconnection networks, [11]), a new project on balanced trees has been launched with Barcelona (work in progress). For the year 2001-2002, C. Banderier will spend a year at MPI to develop random generation methods and LEDA implementations, this under a joint sponsorship of INRIA and MPI.

The INRIA site has produced 16 ALCOM-FT reports [1, 2, 8, 10, 11, 15, 16, 17, 20, 33, 43, 44, 110, 158, 159, 168]. Of these, 9 have been accepted or are appearing in international journals (e.g., *J. of Algorithms*, *Th. Comp. Sc.*), and 5 in international conferences (e.g., *ICALP*).

MPI

The working group Algorithms and Complexity at the Max-Planck-Institute for Computer science (MPI) participates in ALCOM-FT. MPI is involved in Work Packages 1–5. The Algorithms and Complexity group comprises its director, Prof. Kurt Mehlhorn, 15 researcher with Ph.D. (4 with habilitation), and 14 Ph.D. students. More than half of them work partially on ALCOM-FT. Ulrich Meyer is now paid from ALCOM-FT funds.

During the first year of ALCOM-FT, several Ph.D. theses related to the project were completed (most of these are not yet available as ALCOM-FT reports since they are very recent):

Ernst Althaus: Curve Reconstruction and the Traveling Salesman Problem.

Andreas Crauser: LEDA-SM: External Memory Algorithms and Data Structures in Theory and Practice" [73].

Volker Priebe: Average-case complexity of shortest-paths problems.

Mark Ziegelmann: Shortest paths with side constraints and related problems (see also [13]).

Several of our researchers and postdocs come from other ALCOM-FT sites; Berthold Vöcking from Paderborn and Panagiota Fatourou and Spyros Kontogiannis from CTI. Visitors to the site include Susanne Albers, Gemma Casas (Barcelona), Artur Czumaj, Faith Fitch, Alan Frieze, Naveen Garg, Sandeep Sen, and Paul Spirakis (CTI).

Members of the site have participated in most of the important relevant conferences, including ALENEX, APPROX, ESA, FOCS, ICALP, SODA, SPAA, STACS, STOC, SWAT, and WAE. The site hosted several events in September 2000: European Symposium on Algorithms (ESA), Workshop on Algorithm Engineering (WAE), Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX), and Max-Planck Advanced Course on the Foundations of Computer Science.

The 24 ALCOM-FT reports [12, 13, 18, 73, 74, 75, 76, 77, 79, 80, 81, 82, 83, 85, 93, 94, 95, 96, 97, 98, 99, 127, 173, 174] have been authored or co-authored by members of the site. Until now, 22 of these have appeared in conference proceedings or journals, or are scheduled to do so soon.

Paderborn

The activity at site Paderborn is based on the two algorithm groups of Prof. Burkhard Monien and Prof. Friedhelm Meyer auf der Heide at Paderborn University. During the first period, the site is involved in Work Packages 2, 3, 4, and 5 of the ALCOM-FT project. In total, the two groups currently comprise 8 associate/assistant professors and 19 Ph.D. students. Ulf-Peter Schroeder is paid part-time by ALCOM-FT funds.

During the first year of ALCOM-FT, Thomas Decker, Ulf Lorenz, Robert Preis, Marco Riedel, Ingo Rieping, Olaf Schmidt, Ulf-Peter Schroeder, Jürgen Schulze, and Matthias Westermann completed their Ph.D.s with theses in different topics of ALCOM-FT (see www.ub.uni-paderborn.de). Visitors to the site include Rastislav Kralovic and Peter Ruzicka (University of Bratislava), Marios Mavronicolas (University of Cyprus), Michel Toulouse (University of Manitoba), Hal I. Sudborough (University of Dallas), Peter Bro Miltersen (University of Aarhus), Alistair Sinclair (UC Berkeley), and Pankaj Kumar Agarwal (Duke University).

Within the last year Burkhard Monien has been a member of the program committees of MFCS´00, ISPAN´00, and IPDPS´01 (Vice Chair). Moreover, he visited the ALCOM partners in Barcelona. Friedhelm Meyer auf der Heide was involved in the scientific program committees of SPAA´01, ESA´01 (Chair), and ARACNE´01. Members of the site Paderborn have participated in most of the important and relevant conferences, including CPAIOR, ESA, ICALP, ICCS, ISMP, MFCS, SIGGRAPH, SIROCCO, SODA, STACS, and WG. The site hosted the *Second International Workshop of the PAREO working group on Parallel Processing in Operations Research* in June 2000 and the *ALCOM-FT Spring School on Dynamic Algorithms* in May 2001.

The 16 ALCOM-FT reports [48, 69, 70, 71, 86, 87, 88, 89, 90, 91, 100, 116, 132, 133, 134, 169] have been authored or co-authored by members of the site. All of them have appeared in conference proceedings or journals, or are scheduled to do so soon.

CTI

The activity at site Patras is based on the algorithms group at the Computer Technology Institute (Research Unit 1: "Foundations of Computer Science, Relevant Technologies and Applications"). The site is co-ordinating WP 2 and is also involved in WP 1, 2 and 5. The group consists of 8 Faculty member and PhD researchers, and 10 PhD students.

During the first year of ALCOM-FT, the Ph.D Theses of S. Kontogiannis ("Robust Computation in Uncertain Computing Environments") and P. Efraimidis ("Parallel Mixed Integer Linear Programming") have been completed. Two new students (P. Panou and P. Kanellopoulos) have began their PhD. Kurt Mehlhorn from MPI, Alex Shvartsman and Andreas Goerdt have visited Patras.

Members of the site have participated in the Program Committees of major relevant Conferences: ICALP, SPAA, DISC, FCT, EUROPAR, ARACNE etc. P. Spirakis has been invited speaker in Information Days of the Global Computing FET Proactive Initiative. P. Spirakis visited INRIA, Barcelona. CTI has organized back to back in Crete three major conferences: the 33rd ACM STOC, the 13th ACM SPAA and the 28th ICALP.

26 ALCOM-FT reports ([54, 55, 56, 57, 58, 59, 60, 61, 62, 63, 64, 65, 66, 101, 102, 103, 104, 105, 107, 108, 109, 117, 118, 119, 120, 11]) have been authored or co-authored by members of the

site. Until now, 22 of these have appeared in conference proceedings (including STOC, ICALP, SPAA, PODC, MFCS, APPROX, DISC) or journals, or are scheduled to do soon.

Rome

The activity at site Rome is based on the algorithms groups at the Department of Computer and Systems Science of the University of Rome "La Sapienza" and the Department of Computer Science, Systems and Production of the University of Rome "Tor Vergata". The site is involved in Work Packages 1, 2, 4, and 5. In total, the two groups currently consists of 8 faculty (full/associate/assistant professors) and 3 Ph.D. students. Stefano Leonardi, Daniele Frigioni, and Camil Demetrescu are paid part-time by ALCOM-FT funds; Paolo Terrevoli and Silvana Di Vincenzo are assisting with the coordination of the site.

During the first year of ALCOM-FT, Camil Demetrescu has completed his Ph.D. with the thesis *Fully Dynamic Algorithms for Path Problems on Directed Graphs.* Currently, Luigi Laura, Andrea Vitaletti (Rome "La Sapienza") and Fabrizio Grandoni (Rome "Tor Vergata") are working towards their Ph.D. program. Visitors to the site include Susanne Albers, Yossi Azar, Roberto De Prisco, Devdatt Dubhashi, Friedrich Eisenbrand, Thomas Erlebach, Paolo Ferragina, Amos Fiat, Peter Gacs, Nicola Galesi, Jens Lagergren, Kurt Mehlhorn, Franco P. Preparata, Kirk Pruhs, Leen Stougie, R. Ravi, Alberto Sangiovanni Vincentelli, Luca Trevisan, Dorothea Wagner, and Emo Welzl.

Members of the site have participated in most of the important relevant conferences, including ALENEX, APPROX, ESA, FOCS, FST&TCS, MFCS, SODA, STOC, VL, and WAE. The site organizes a permanent seminar on algorithms (SIA) in cooperation with the Department of Computer Science of the University of Roma "La Sapienza". See: `http://www.dis.uniroma1.it/~algo`. The site will host an upcoming Summer School on Algorithm Engineering, to be held on September 10–12, 2001, `www.info.uniroma2.it/~italiano/School` and will host the first Project Workshop and Review Meeting.

The 23 ALCOM-FT reports [19, 22, 23, 24, 25, 26, 27, 28, 29, 30, 31, 32, 40, 41, 42, 45, 46, 47, 102, 147, 148, 149, 150] have been authored or co-authored by members of the site. Until now, 18 of these have appeared in conference proceedings or journals, or are scheduled to do so soon.

Utrecht

The people participating in ALCOM-FT at at site Utrecht are members of the group Algorithm Design, which is part of the Department of Information and Computing Sciences at Utrecht University. The group participates in WP 3 and 4. The group currently consists of 7 faculty members/researchers and, as of July 1, 3 Ph.D. students. The attempts to hire a Ph.D. student and a postdoc fully paid by ALCOM-FT have not been successful yet.

During the first year of ALCOM-FT, Bram Verweij completed his Ph.D. with the thesis *Selected Applications of Integer Programming: A Computational Study.* Marc Uetz has visited Utrecht for one week.

Members of the site have participated in a number of important relevant conferences, including APPROX, ESA, ICALP, IPCO, MAPSP, SOFSEM and WG. The conference *Integer Programming and Combinatorial Optimization*, co-organized by Karen Aardal, was held in Utrecht in June 2001.

Members of the site have authored or co-authored the ALCOM-FT reports [106, 111, 112, 113, 114, 115, 121, 122, 125, 129, 160, 161, 162, 163]. Several of these reports have appeared in conference proceedings or journals, or are scheduled to do so soon.

Warwick

The ALCOM-FT participants at the Warwick site are the Algorithms Group of the Department of Computer Science at the University of Warwick. Warwick is involved in Work Packages 1,2, and

4. The Group currently consists of five permanent academic staff (Hristo Djidjev, Leslie Goldberg, Paul Goldberg, Mike Paterson and Alex Tiskin), two research fellows (Petra Berenbrink and Tom Friedetzky), and three PhD students. Friedetzky is paid by ALCOM-FT. He and Tiskin assist with the coordination of the project.

During the year, Hesham Al-Ammal completed his PhD thesis on *Stability and Performance of Contention Resolution Protocols* under the supervision of Leslie Goldberg. Visitors to Warwick included Richard Cole, Persi Diaconis, Mark Jerrum, and Don Knuth.

In May 2001 Warwick organized *Warwick Algorithms Day* at which the speakers were Richard Cole (Courant Institute, NYU), Wilfrid Kendall (Statistics, Warwick), Martin Dyer (Leeds), Paul Goldberg (Warwick), Graham Brightwell (London School of Economics), Mark Jerrum (Edinburgh), Leszek Gasienic (Liverpool) and David Manlove (Glasgow).

Also in May, Warwick hosted this year's Hardy Lecturer, Persi Diaconis (Stanford), who talked about the Metropolis Algorithm.

ALCOM affiliated researchers at Warwick have participated in most of the important relevant conferences, including ALENEX, APPROX, COCOON, ESA, FOCS, ICALP, SPAA, STOC, and SWAT, and attended workshops at Dagstuhl and Oberwolfach.

Eleven ALCOM-FT reports [3, 4, 5, 6, 7, 9, 14, 72, 146, 153, 154] have been authored or co-authored by members of the site. Most of these have appeared in conference proceedings or journals, or are scheduled to do so soon.

Part III Scientific Reports

In the following, we list all scientific reports which has been published in the ALCOM-FT Technical Report Series in the period June 1, 2000 to May 31, 2001. All reports are available online at `www.brics.dk/ALCOM-FT/TR`.

References

[1] Philippe Flajolet and Brigitte Vallée. Continued Fractions, Comparison Algorithms, and Fine Structure Constants. Technical Report ALCOMFT-TR-01-1, ALCOM-FT, INRIA, January 2001. APPEARANCE: In Michel Théra, editor, *Constructive, Experimental, and Nonlinear Analysis*, volume 27 of *Canadian Mathematical Society Conference Proceedings*, pages 53–82, Providence, 2000. American Mathematical Society.

[2] Philippe Flajolet, Xavier Gourdon, and Daniel Panario. The complete analysis of a polynomial factorization algorithm over finite fields. Technical Report ALCOMFT-TR-01-2, ALCOM-FT, INRIA, January 2001.

[3] Alexandre Tiskin. A new way to divide and conquer. Technical Report ALCOMFT-TR-01-3, ALCOM-FT, Warwick, January 2001.

[4] Leslie Ann Goldberg. Computation in permutation groups: counting and randomly sampling orbits. Technical Report ALCOMFT-TR-01-4, ALCOM-FT, Warwick, January 2001. APPEARANCE: In J. Hirschfeld, editor, *Surveys in Combinatorics*. Cambridge University Press, 2001. To appear.

[5] Martin Dyer, Leslie Ann Goldberg, Catherine Greenhill, Gabriel Istrate, and Mark Jerrum. Convergence of the Iterated Prisoner's Dilemma Game. Technical Report ALCOMFT-TR-01-5, ALCOM-FT, Warwick, January 2001.

[6] Martin Dyer, Leslie Ann Goldberg, Catherine Greenhill, and Mark Jerrum. On the relative complexity of approximate counting problems. Technical Report ALCOMFT-TR-01-6, ALCOM-FT, Warwick, January 2001. APPEARANCE: In *Proceedings of APPROX 2000*, volume 1913 of *Lecture Notes in Computer Science*, pages 108–119. Springer-Verlag, 2000.

[7] L.A. Goldberg, M. Jerrum, S. Kannan, and M. Paterson. A Bound on the Capacity of Backoff and Acknowledgement-Based Protocols. Technical Report ALCOMFT-TR-01-7, ALCOM-FT, Warwick, January 2001. APPEARANCE: In U. Montanari, J. Rolin, and E. Welzl, editors, *Proceedings of ICALP 2000*, volume 1853 of *Lecture Notes in Computer Science*, pages 705–716. Springer-Verlag, 2000.

[8] Philippe Flajolet and Robert Sedgewick. Analytic Combinatorics: Functional Equations, Rational and Algebraic Functions. Technical Report ALCOMFT-TR-01-8, ALCOM-FT, INRIA, January 2001.

[9] M. Adler, F. Fich, L. A. Goldberg, and M. Paterson. Tight Size Bounds for Packet Headers in Narrow Meshes. Technical Report ALCOMFT-TR-01-9, ALCOM-FT, Warwick, January 2001. APPEARANCE: In U. Montanari, J. Rolin, and E. Welzl, editors, *Proceedings of ICALP 2000*, volume 1853 of *Lecture Notes in Computer Science*, pages 756–767. Springer-Verlag, 2000.

[10] Jean-Francois Dantzer and Philippe Robert. Analysis of a multi-class queueing system. Technical Report ALCOMFT-TR-01-10, ALCOM-FT, INRIA, January 2001.

[11] Philippe Flajolet, Kostas Hatzis, Sotiris Nikoletseas, and Paul Spirakis. On the Robustness of Interconnections in Random Graphs:A Symbolic Approach. Technical Report ALCOMFT-TR-01-11, ALCOM-FT, INRIA and Patras, January 2001.

[12] P. Sanders and R. Solis-Oba. How Helpers Hasten h-Relations. Technical Report ALCOMFT-TR-01-12, ALCOM-FT, MPI, January 2001. APPEARANCE: In *8th European Symposium on Algorithms (ESA)*, number 1879 in LNCS, pages 392–402. Springer, 2000.

[13] Kurt Mehlhorn and Mark Ziegelmann. CNOP - A Package for Constrained Network Optimization. Technical Report ALCOMFT-TR-01-13, ALCOM-FT, MPI, February 2001. APPEARANCE: In *3rd Workshop on Algorithm Engineering and Experiments (ALENEX01)*, LNCS, 2001.

[14] Alexandre Tiskin. Parallel convex hull computation by generalised regular sampling. Technical Report ALCOMFT-TR-01-14, ALCOM-FT, Warwick, February 2001.

[15] Philippe Flajolet and Guy Louchard. Analytic Variations on the Airy Distribution. Technical Report ALCOMFT-TR-01-15, ALCOM-FT, INRIA, February 2001.

[16] Alain Dupuis, Fabrice Guillemin, and Philippe Robert. Modeling Max-Min Fairness for Elastic Flows in Telecommunication Networks. Technical Report ALCOMFT-TR-01-16, ALCOM-FT, INRIA, February 2001. APPEARANCE: In *ITC'17*, Salvador da Bahia, September 2001.

[17] C. Banderier, M. Bousquet-Mélou, A. Denise, P. Flajolet, D. Gardy, and D. Gouyou-Beauchamps. Generating functions for generating trees. Technical Report ALCOMFT-TR-01-17, ALCOM-FT, INRIA, February 2001.

[18] Kurt Mehlhorn and Guido Schäfer. A Heuristic for Dijkstra's Algorithm with Many Targets and its Use in Weighted Matching Algorithms. Technical Report ALCOMFT-TR-01-18, ALCOM-FT, MPI, February 2001. APPEARANCE: In *Proceedings of the 9th Annual European Symposium on Algorithms (ESA 2001)*, Lecture Notes in Computer Science, to appear. Springer, 2001.

[19] Giorgio Ausiello, Paolo G. Franciosa, Daniele Frigioni, and Roberto Giaccio. Decremental Maintenance of Minimum Rank Hyperpaths and Minimum Models of Horn Formulae. Technical Report ALCOMFT-TR-01-19, ALCOM-FT, Rome, February 2001.

[20] Cyril Banderier, Philippe Flajolet, Gilles Schaeffer, and Michele Soria. Random Maps, Coalescing Saddles, Singularity Analysis, and Airy Phenomena. Technical Report ALCOMFT-TR-01-20, ALCOM-FT, INRIA, March 2001.

[21] José Luis Balcázar, Jorge Castro, and David Guijarro. A General Dimension for Exact Learning. Technical Report ALCOMFT-TR-01-21, ALCOM-FT, Barcelona, March 2001.

[22] C. Demetrescu, I. Finocchi, and G. Liotta. Visualizing Algorithms over the Web with the Publication-driven Approach. Technical Report ALCOMFT-TR-01-22, ALCOM-FT, Rome, March 2001. APPEARANCE: In *Proceedings of the 4th Workshop on Algorithm Engineering (WAE'00)*, 2000.

[23] C. Demetrescu, D. Frigioni, A. Marchetti-Spaccamela, and U. Nanni. Maintaining Shortest Paths in Digraphs with Arbitrary Arc Weights: An Experimental Study. Technical Report ALCOMFT-TR-01-23, ALCOM-FT, Rome, March 2001. APPEARANCE: In *Proceedings of the 4th Workshop on Algorithm Engineering (WAE'00)*, 2000.

[24] C. Demetrescu and G.F. Italiano. What Do We Learn from Experimental Algorithmics? Technical Report ALCOMFT-TR-01-24, ALCOM-FT, Rome, March 2001. APPEARANCE: In *Proceedings of the 25th International Symposium on Mathematical Foundations of Computer Science (MFCS'00)*, 2000.

[25] C. Demetrescu and G.F. Italiano. Fully Dynamic Transitive Closure: Breaking Through the $O(n^2)$ Barrier. Technical Report ALCOMFT-TR-01-25, ALCOM-FT, Rome, March 2001. Appearance: In *Proc. 41st IEEE Symp. on Foundations of Computer Science (FOCS00)*, 2000.

[26] C. Demetrescu and I. Finocchi. Smooth Animation of Algorithms in a Declarative Framework. Technical Report ALCOMFT-TR-01-26, ALCOM-FT, Rome, March 2001. Appearance: *Special Issue of the Journal of Visual Languages and Computing devoted to selected papers of the 15-th IEEE International Symposium on Visual Languages (VL99).* To appear.

[27] Massimiliano Caramia, Paolo Dell'Olmo, and Giuseppe F. Italiano. New Algorithms for Examination Timetabling. Technical Report ALCOMFT-TR-01-27, ALCOM-FT, Rome, March 2001. Appearance: In *Proc. 4th Workshop on Algorithm Engineering (WAE 2000)*, 2000.

[28] Massimiliano Caramia, Paolo Dell'Olmo, and Giuseppe F. Italiano. CHECKCOL: Improved Local Search for Graph Coloring. Technical Report ALCOMFT-TR-01-28, ALCOM-FT, Rome, March 2001.

[29] Alberto Caprara, Giuseppe F. Italiano, G. Mohan, Alessandro Panconesi, and Aravind Srinivasan. Wavelength Rerouting in Optical Networks, or the Venetian Routing Problem. Technical Report ALCOMFT-TR-01-29, ALCOM-FT, Rome, March 2001. Appearance: In *Proc. 3rd International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX 2000)*, 2000.

[30] Michele Flammini and Gaia Nicosia. On Multicriteria Online Problems. Technical Report ALCOMFT-TR-01-30, ALCOM-FT, Rome, March 2001. Appearance: In *Proceedings of the 8th Annual European Symposium on Algorithms (ESA)*, volume 1879 of *Lecture Notes in Computer Science*, pages 191–201. Springer-Verlag, September 2000.

[31] Daniele Frigioni, Alberto Marchetti-Spaccamela, and Umberto Nanni. Fully Dynamic Shortest Paths in Digraphs with Arbitrary Arc Weights. Technical Report ALCOMFT-TR-01-31, ALCOM-FT, Rome, March 2001. Appearance: *Journal of Algorithms*. To appear.

[32] Serafino Cicerone, Gabriele Di Stefano, Daniele Frigioni, and Umberto Nanni. A Fully Dynamic Algorithm for Distributed Shortest Paths. Technical Report ALCOMFT-TR-01-32, ALCOM-FT, Rome, March 2001. Appearance: *Theoretical Computer Science, to appear.*

[33] Vincent Dumas, Fabrice Guillemin, and Philippe Robert. Limit results for Markovian models of TCP. Technical Report ALCOMFT-TR-01-33, ALCOM-FT, INRIA, April 2001.

[34] Gerth Stølting Brodal and Riko Jacob. Dynamic Planar Convex Hull with Optimal Query Time and $O(\log n \log\log n)$ Update Time. Technical Report ALCOMFT-TR-01-34, ALCOM-FT, Aarhus, April 2001. Appearance: In *Proc. 7th Scandinavian Workshop on Algorithm Theory*, volume 1851 of *Lecture Notes in Computer Science*, pages 57–70. Springer Verlag, Berlin, 2000.

[35] Stephen Alstrup, Gerth Stølting Brodal, and Theis Rauhe. New Data Structures for Orthogonal Range Searching. Technical Report ALCOMFT-TR-01-35, ALCOM-FT, Aarhus, April 2001. Appearance: In *Proc. 41st Annual Symposium on Foundations of Computer Science*, pages 198–207, 2000.

[36] J.A. Makowsky and K. Meer. Polynomials of bounded tree-width. Technical Report ALCOMFT-TR-01-36, ALCOM-FT, Aarhus, April 2001.

[37] K. Meer and G.W. Weber. Some aspects of studying an optimization or decision problem in different computational models. Technical Report ALCOMFT-TR-01-37, ALCOM-FT, Aarhus, April 2001.

[38] Lars Arge, Gerth Stølting Brodal, and Laura Toma. On External-Memory MST, SSSP and Multi-way Planar Graph Separation. Technical Report ALCOMFT-TR-01-38, ALCOM-FT, Aarhus, April 2001. APPEARANCE: In *Proc. 7th Scandinavian Workshop on Algorithm Theory*, volume 1851 of *Lecture Notes in Computer Science*, pages 433–447. Springer Verlag, Berlin, 2000.

[39] J. Baixeries, G. Casas-Garriga, and J.L. Balcázar. A Faster Algorithm for Finding Frequent Sets. Technical Report ALCOMFT-TR-01-39, ALCOM-FT, Barcelona, April 2001.

[40] Giuseppe Cattaneo, Umberto Ferraro-Petrillo, and Giuseppe F. Italiano. CATAI: Concurrent Algorithms and data Types Animationover the Internet. Technical Report ALCOMFT-TR-01-40, ALCOM-FT, Rome, April 2001.

[41] Massimiliano Curcio, Stefano Leonardi, and Andrea Vitaletti. Integrated Prefetching and Caching for the World Wide Web via User Cooperation. Technical Report ALCOMFT-TR-01-41, ALCOM-FT, Rome, April 2001.

[42] Luca Becchetti, Stefano Leonardi, Alberto Marchetti-Spaccamela, and Kirk Pruhs. Online Weighted Flow Time and Deadline Scheduling. Technical Report ALCOMFT-TR-01-42, ALCOM-FT, Rome, April 2001.

[43] Philippe Flajolet, Yves Guivarc'h, Wojciech Szpankowski, and Brigitte Vallée. Hidden Pattern Statistics. Technical Report ALCOMFT-TR-01-43, ALCOM-FT, INRIA, April 2001.

[44] Christine Fricker, Philippe Robert, and Danielle Tibi. On the fluid limits of some loss networks. Technical Report ALCOMFT-TR-01-44, ALCOM-FT, INRIA, April 2001.

[45] L. Becchetti, S. Leonardi, and M. Muthukrishnan. Average Stretch without Migration. Technical Report ALCOMFT-TR-01-45, ALCOM-FT, Rome, April 2001.

[46] Luca Becchetti and Stefano Leonardi. Non-Clairvoyant Scheduling to Minimize the Average Flow Time on Single and Parallel Machines. Technical Report ALCOMFT-TR-01-46, ALCOM-FT, Rome, April 2001.

[47] Stefano Leonardi, Alberto Marchetti-Spaccamela, and andrea Vitaletti. Approximation algorithms for bandwidth and storage allocation under real time constraints. Technical Report ALCOMFT-TR-01-47, ALCOM-FT, Rome, April 2001. APPEARANCE: In *Proc. of FST TCS 00, LNCS 1974*, pages 409–420. Springer-Verlag, 2000.

[48] Friedhelm Meyer auf der Heide, Harald Räcke, and Matthias Westermann. Data Management in Hierarchical Bus Networks. Technical Report ALCOMFT-TR-01-48, ALCOM-FT, Paderborn, April 2001. APPEARANCE: In *Proceedings of the 12th ACM Symposium on Parallel Algorithms and Architectures (SPAA)*, pages 109–118, 2000.

[49] Eric Bach, Joan Boyar, Leah Epstein, Lene M. Favrholdt, Tao Jiang, Kim S. Larsen, Guo-Hui Lin, and Rob van Stee. Tight Bounds on the Competitive Ratio on Accommodating Sequences. Technical Report ALCOMFT-TR-01-49, ALCOM-FT, Aarhus, May 2001. APPEARANCE: *Journal of Scheduling*. To appear.

[50] Yossi Azar, Joan Boyar, Leah Epstein, Lene M. Favrholdt, Kim S. Larsen, and Morten N. Nielsen. Fair versus Unrestricted Bin Packing. Technical Report ALCOMFT-TR-01-50, ALCOM-FT, Aarhus, May 2001. APPEARANCE: In *Proceedings of the Seventh Scandinavian Workshop on Algorithm Theory*, volume 1851 of *Lecture Notes in Computer Science*, pages 200–213. Springer-Verlag, 2000. (Preliminary version).

[51] Kim S. Larsen. Relaxed Multi-Way Trees with Group Updates. Technical Report ALCOMFT-TR-01-51, ALCOM-FT, Aarhus, May 2001. APPEARANCE: In *Proceedings of the Twentieth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, pages 93–101. ACM Press, 2001.

[52] Lars Jacobsen and Kim S. Larsen. Variants of (a,b)-Trees with Relaxed Balance. Technical Report ALCOMFT-TR-01-52, ALCOM-FT, Aarhus, May 2001. APPEARANCE: *International Journal of Foundations of Computer Science*. To appear.

[53] Stephen Alstrup, Gerth Stølting Brodal, and Theis Rauhe. Optimal Static Range Reporting in One Dimension. Technical Report ALCOMFT-TR-01-53, ALCOM-FT, Aarhus, May 2001. APPEARANCE: In *Proc. 33rd Annual ACM Symposium on Theory of Computing*, 2001.

[54] Ioannis Caragiannis, Afonso Ferreira, Christos Kaklamanis, Stephane Perennes, and Herve Rivano. Fractional Path Coloring with Applications to WDM Networks. Technical Report ALCOMFT-TR-01-54, ALCOM-FT, Patras, May 2001. APPEARANCE: In *28th International Colloquium on Automata, Languages and Programming (ICALP '01)*. To appear.

[55] Ioannis Caragiannis, Christos Kaklamanis, and Pino Persiano. Wavelength Routing in All–Optical Tree Networks: A Survey. Technical Report ALCOMFT-TR-01-55, ALCOM-FT, Patras, May 2001.

[56] Ioannis Caragiannis, Christos Kaklamanis, and Ioannis Vergados. Greedy Dynamic Hot-Potato Routing on Arrays. Technical Report ALCOMFT-TR-01-56, ALCOM-FT, Patras, May 2001. APPEARANCE: In *2000 International Symposium on Parallel Algorithms and Networks (I-SPAN 2000)*, pages 178–185. IEEE Computer Society Press, 2000.

[57] Ioannis Caragiannis, Christos Kaklamanis, and Evi Papaioannou. Competitive Analysis of On–line Randomized Call Control in Cellular Networks. Technical Report ALCOMFT-TR-01-57, ALCOM-FT, Patras, May 2001. APPEARANCE: In *15th International Parallel and Distributed Processing Symposium (IPDPS '01)*. IEEE Computer Society Press, 2001.

[58] Constantinos Bartzis, Ioannis Caragiannis, Christos Kaklamanis, and Ioannis Vergados. Experimental Evaluation of Hot-Potato Routing Algorithms in 2-Dimensional Processor Arrays. Technical Report ALCOMFT-TR-01-58, ALCOM-FT, Patras, May 2001. APPEARANCE: In *EUROPAR-Parallel Processing 2000*, number 1900 in LNCS, pages 877–881. Springer, 2000.

[59] Vincenzo Auletta, Ioannis Caragiannis, Christos Kaklamanis, and Pino Persiano. Randomized Path Coloring on Binary Trees. Technical Report ALCOMFT-TR-01-59, ALCOM-FT, Patras, May 2001. APPEARANCE: In *Proceedings of the 3rd Workshop on Approximation Algorithms for Combinatorial Optimization Algorithm Engineering – APPROX 2000*. Lect. Notes in Comp. Science, Springer-Verlag, 2000.

[60] Ioannis Caragiannis, Christos Kaklamanis, and Evi Papaioannou. Efficient On-line Communication in Cellular Networks. Technical Report ALCOMFT-TR-01-60, ALCOM-FT, Patras, May 2001. APPEARANCE: In *Proceedings of the 12th Symposium on Parallel Algorithms and Architectures – SPAA 2000*, pages 46–53. ACM, 2000.

[61] Ioannis Chatzigiannakis, Sotiris Nikoletseas, and Paul Spirakis. Analysis and Experimental Evaluation of an Innovative and Efficient Routing Protocol for Ad-hoc Mobile Networks. Technical Report ALCOMFT-TR-01-61, ALCOM-FT, Patras, May 2001. APPEARANCE: In *Algorithm Engineering – WAE 2000*. Lect. Notes in Comp. Science, Springer-Verlag, 2000.

[62] Ioannis Chatzigiannakis, Sotiris Nikoletseas, and Paul Spirakis. An Efficient Routing Protocol for Hierarchical Ad-hoc Mobile Networks. Technical Report ALCOMFT-TR-01-62, ALCOM-FT, Patras, May 2001. APPEARANCE: In *Proceedings of the IPDPS Workshop on Parallel and Distributed Computing Issues in Wireless Networks and Mobile Computing – PDP 2001*. IEEE, 2001.

[63] Sotiris Nikoletseas, Krishna Palem, Paul Spirakis, and Moti Yung. Connectivity Properties in Random Regular Graphs with Edge Faults. Technical Report ALCOMFT-TR-01-63, ALCOM-FT, Patras, May 2001. APPEARANCE: *International Journal of Foundations of Computer Science – IJFCS*, pages 247–262, 2000.

[64] John Garofalakis, Sotiris Nikoletseas, and Paul Spirakis. Latency Lower Bounds for Randomized Dynamic Bandwidth Allocation. Technical Report ALCOMFT-TR-01-64, ALCOM-FT, Patras, May 2001. APPEARANCE: In *Proceedings of the 14th Symposium on Distributed Computing – DISC 2000*, 2000.

[65] Sotiris Nikoletseas and Paul Spirakis. Efficient Communication Establishment in Adverse Communication Environments. Technical Report ALCOMFT-TR-01-65, ALCOM-FT, Patras, May 2001. APPEARANCE: In *Proceeding of the ICALP Satellite Workshop on Approximation and Randomized Algorithms in Communication Networks – ARACNE 2000*. EATCS, 2000.

[66] Sotiris Nikoletseas and Paul Spirakis. Randomized Techniques for Modelling Faults and Achieving Robust Computing. Technical Report ALCOMFT-TR-01-66, ALCOM-FT, Patras, May 2001. APPEARANCE: In *Handbook of Randomized Computing*. Kluwer Academic Publishers, 2001. To appear.

[67] Rasmus Pagh. On the Cell Probe Complexity of Membership and Perfect Hashing. Technical Report ALCOMFT-TR-01-67, ALCOM-FT, Aarhus, May 2001. APPEARANCE: In *Proc. 33rd Annual ACM Symposium on Theory of Computing*, 2001.

[68] Rasmus Pagh. A Trade-Off For Worst-Case Efficient Dictionaries. Technical Report ALCOMFT-TR-01-68, ALCOM-FT, Aarhus, May 2001. APPEARANCE: *Nordic J. Comput.*, 7(3):151–163, 2000.

[69] André Brinkmann, Kay Salzwedel, and Christian Scheideler. Efficient, Distributed Data Placement Strategies for Storage Area Networks. Technical Report ALCOMFT-TR-01-69, ALCOM-FT, Paderborn, May 2001. APPEARANCE: In *Proceedings of the 12th Annual ACM Symposium on Parallel Algorithms and Architectures*, pages 119–128, 2000.

[70] Ben H.H. Juurlink, Petr Kolman, Friedhelm Meyer auf der Heide, and Ingo Rieping. Optimal Broadcast on Parallel Locality Models. Technical Report ALCOMFT-TR-01-70, ALCOM-FT, Paderborn, May 2001. APPEARANCE: In *Proceedings of the 7th Colloquium on Structural Information and Communication Complexity (SIROCCO)*, pages 211–225, 2000.

[71] Christof Krick, Harald Räcke, and Matthias Westermann. Approximation Algorithms for Data Management in Networks. Technical Report ALCOMFT-TR-01-71, ALCOM-FT, Paderborn, May 2001. APPEARANCE: In *Proceedings of the 13th ACM Symposium on Parallel Algorithms and Architectures (SPAA)*, 2001. To appear.

[72] Petra Berenbrink, Tom Friedetzky, and Leslie Ann Goldberg. The Natural Work-Stealing Algorithm is Stable. Technical Report ALCOMFT-TR-01-72, ALCOM-FT, Warwick, May 2001.

[73] Andreas Crauser. LEDA-SM:External Memory Algorithms and Data Structures in Theory and Practice. Technical Report ALCOMFT-TR-01-73, ALCOM-FT, MPI, May 2001. APPEARANCE: PhD thesis, Universität des Saarlandes and MPII, 2001.

[74] Stefan Burkhardt and Juha Kärkkäinen. Better Filtering with Gapped q-Grams. Technical Report ALCOMFT-TR-01-74, ALCOM-FT, MPI, May 2001. APPEARANCE: In *Proc. 12th Annual Symposium on Combinatorial Pattern Matching*, 2001. To appear.

[75] Ernst Althaus, Denys Duchier, Alexander Koller, Kurt Mehlhorn, Joachim Niehren, and Sven Thiel. An Efficient Algorithm for the Configuration Problem of Dominance Graphs. Technical Report ALCOMFT-TR-01-75, ALCOM-FT, MPI, May 2001. APPEARANCE: In *Proceedings of the 12th ACM-SIAM Symposium on Discrete Algoritms (SODA 01)*, pages 815–824, Washington DC, USA, January 2001. ACM-SIAM, ACM-SIAM.

[76] Panagiota Fatourou and Maurice Herlihy. Brief Announcement: Adding Networks. Technical Report ALCOMFT-TR-01-76, ALCOM-FT, MPI, May 2001. APPEARANCE: In *Proceedings of the 20th ACM Symposium on Principles of Distributed Computing*, Newport, Rhode Island, August 2001.

[77] Panagiota Fatourou. Low-Contention Depth-First Scheduling of Parallel Computations with Write-Once Synchronization Variables. Technical Report ALCOMFT-TR-01-77, ALCOM-FT, MPI, May 2001. APPEARANCE: In *Proceedings of the 13th ACM Symposium on Parallel Algorithms and Architectures*, Crete Island, Greece, July 2001.

[78] Lars Jacobsen, Kim S. Larsen, and Morten N. Nielsen. On the Existence and Construction of Non-Extreme (a,b)-Trees. Technical Report ALCOMFT-TR-01-78, ALCOM-FT, Aarhus, May 2001.

[79] David A. Hutchinson, Peter Sanders, and Jeffrey S. Vitter. Duality between prefetching and queued writing with Parallel Disks. Technical Report ALCOMFT-TR-01-79, ALCOM-FT, MPI, May 2001. APPEARANCE: In *9th European Symposium on Algorithms (ESA)*, LNCS. Springer, 2001. To appear.

[80] Catherine Mc Geoch, Peter Sanders, Rudolf Fleischer, Paul R. Cohen, and Doina Precup. Searching for Big-Oh in the Data: Inferring Asymptotic Complexity from Experiments. Technical Report ALCOMFT-TR-01-80, ALCOM-FT, MPI, May 2001. APPEARANCE: In *Workshop on Algorithm Engineering*, LNCS, 2000. to appear.

[81] Peter Sanders. Presenting Data from Experiments in Algorithmics. Technical Report ALCOMFT-TR-01-81, ALCOM-FT, MPI, May 2001.

[82] Peter Sanders. Reconciling Simplicity and Realism in Parallel Disk Models. Technical Report ALCOMFT-TR-01-82, ALCOM-FT, MPI, May 2001. APPEARANCE: In *12th ACM-SIAM Symposium on Discrete Algorithms*, pages 67–76, Washington DC, 2001.

[83] Peter Sanders. Fast Priority Queues for Cached Memory. Technical Report ALCOMFT-TR-01-83, ALCOM-FT, MPI, May 2001. APPEARANCE: *ACM Journal of Experimental Algorithmics*, 5, 2000.

[84] Matthias Elf, Michael Jünger, and Giovanni Rinaldi. Minimizing Breaks by Maximizing Cuts. Technical Report ALCOMFT-TR-01-84, ALCOM-FT, Cologne, May 2001.

[85] Spyros Kontogiannis, Grammati Pantziou, Paul Spirakis, and Moti Yung. Robust parallel computations through randomization. Technical Report ALCOMFT-TR-01-85, ALCOM-FT, Patras and MPI, May 2001. APPEARANCE: *Theory of Computing Systems*, 33:427–464, 2000.

[86] Robert Elsässer, Rastislav Královic, and Burkhard Monien. Scalable Sparse Topologies with Small Spectrum. Technical Report ALCOMFT-TR-01-86, ALCOM-FT, Paderborn, May 2001. APPEARANCE: In *Proceedings of the 18th International Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 218–229, 2001.

[87] Meinolf Sellmann and Torsten Fahle. Coupling Variable Fixing Algorithms for the Automatic Recording Problem. Technical Report ALCOMFT-TR-01-87, ALCOM-FT, Paderborn, May 2001. APPEARANCE: In *Proceedings of the 9th Annual European Symposium on Algorithms (ESA)*, 2001. To appear.

[88] Norbert Sensen. Lower Bounds and Exact Algorithms for the Graph Partitioning Problem using Multicommodity Flows. Technical Report ALCOMFT-TR-01-88, ALCOM-FT, Paderborn, May 2001. Appearance: In *Proceedings of the 9th Annual European Symposium on Algorithms (ESA)*, 2001. To appear.

[89] Robert Elsässer, Thomas Lücking, and Burkhard Monien. New Spectral Bounds on k-Partitioning of Graphs. Technical Report ALCOMFT-TR-01-89, ALCOM-FT, Paderborn, May 2001. Appearance: In *Proceedings of the 13th ACM Symposium on Parallel Algorithms and Architectures (SPAA)*, 2001. To appear.

[90] Meinolf Sellmann, Kyriakos Zervoudakis, Panagiotis Stamatopoulos, and Torsten Fahle. Crew Assignment via Constraint Programming: Integrating Column Generation and Heuristic Tree Search. Technical Report ALCOMFT-TR-01-90, ALCOM-FT, Paderborn, May 2001. Appearance: In *Proceedings of the 2nd Intern. Workhop on Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems, (CP-AI-OR'00)*, 2000. Extended version, submitted.

[91] Torsten Fahle, Ulrich Junker, Stefan E. Karisch, Niklas Kohl, Meinolf Sellmann, and Bo Vaaben. Constraint Programming Based Column Generation for Crew Assignment. Technical Report ALCOMFT-TR-01-91, ALCOM-FT, Paderborn, May 2001. Appearance: *to appear in Journal of Heuristics*, 8(1), 2002. Earlier version presented at CP-AI-OR'99.

[92] José L. Balcázar, Yang Dai, and Osamu Watanabe. Random sampling techniques for training support vector machines. Technical Report ALCOMFT-TR-01-92, ALCOM-FT, Barcelona, May 2001.

[93] Lars Arge, Ulrich Meyer, Laura Toma, and Norbert Zeh. On External-Memory Planar Depth First Search. Technical Report ALCOMFT-TR-01-93, ALCOM-FT, MPI, May 2001. Appearance: In *Proc. International Workshop on Algorithms and Data Structures (WADS)*, LNCS. Springer, to be published 2001.

[94] Ulrich Meyer. Heaps are better than Buckets: Parallel Shortest Paths on unbalanced Graphs. Technical Report ALCOMFT-TR-01-94, ALCOM-FT, MPI, May 2001. Appearance: In *Proc. Euro-Par 2001 Parallel Processing*, LNCS. Springer, to be published 2001.

[95] Ulrich Meyer and Peter Sanders. Parallel Shortest Path for Arbitrary Graphs. Technical Report ALCOMFT-TR-01-95, ALCOM-FT, MPI, May 2001. Appearance: In *Proc. Euro-Par 2000 Parallel Processing*, volume 1900 of *LNCS*, pages 461–470. Springer, 2000.

[96] Ulrich Meyer. External Memory BFS on Undirected Graphs with Bounded Degree. Technical Report ALCOMFT-TR-01-96, ALCOM-FT, MPI, May 2001. Appearance: In *Proc. 12th Annual Symposium on Discrete Algorithms*, pages 87–88. ACM–SIAM, 2001.

[97] Rajmohan Rajaraman, Andréa W. Richa, Berthold Vöcking, and Gayathri Vuppuluri. A Data Tracking Scheme for General Networks. Technical Report ALCOMFT-TR-01-97, ALCOM-FT, MPI, May 2001. Appearance: In *ACM Symposium on Parallel Algorithms and Architectures*, pages 247–254, 2001.

[98] Ulrich Meyer. Directed Single–Source Shortest–Paths in Linear Average–Case Time. Technical Report ALCOMFT-TR-01-98, ALCOM-FT, MPI, May 2001. Appearance: In *Proc. 12th Annual Symposium on Discrete Algorithms*, pages 797–806. ACM–SIAM, 2001.

[99] Peter Sanders. Asynchronous Scheduling of Redundant Disk Arrays. Technical Report ALCOMFT-TR-01-99, ALCOM-FT, MPI, May 2001. Appearance: In *12th ACM Symposium on Parallel Algorithms and Architectures*, pages 89–98, 2000.

[100] Meinolf Sellmann and Torsten Fahle. CP-based Lagrangian Relaxation for a Multimedia Application. Technical Report ALCOMFT-TR-01-100, ALCOM-FT, Paderborn, May 2001. Appearance: In *Proceedings of the 3rd Intern. Workhop on Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems, (CP-AI-OR'01)*, 2001.

[101] Sotiris Nikoletseas, Grigorios Prasinos, Paul Spirakis, and Christos Zaroliagis. Attack Propagation in Networks. Technical Report ALCOMFT-TR-01-101, ALCOM-FT, Patras, May 2001. Appearance: In *Proc. 13th ACM Symposium on Parallel Algorithms and Architectures (SPAA 2001)*. ACM Press, 2001. To appear.

[102] Daniele Frigioni, Tobias Miller, Umberto Nanni, and Christos Zaroliagis. An Experimental Study of Dynamic Algorithms for Transitive Closure. Technical Report ALCOMFT-TR-01-102, ALCOM-FT, Patras and Rome, May 2001.

[103] Alexis Kaporis, Lefteris Kirousis, Yannis Stamatiou, Malvina Vamvakari, and Michele Zito. The unsatisfiability threshold revisited. Technical Report ALCOMFT-TR-01-103, ALCOM-FT, Patras, May 2001. Appearance: In *LICS Workshop on Satisfiability Testing: Theory and Applications – SAT 2001*, 2001.

[104] Alexis C. Kaporis, Lefteris M. Kirousis, and Yannis C. Stamatiou. A note on the non-colorability threshold of a random graph. Technical Report ALCOMFT-TR-01-104, ALCOM-FT, Patras, May 2001. Appearance: *Electronic Journal of Combinatorics*, 2001.

[105] Lefteris M. Kirousis, Evangelos Kranakis, Danny Krizanc, and Yannis C. Stamatiou. Locating Information with Uncertainty in Fully Interconnected Networks: The Case of Non-Distributed Memory. Technical Report ALCOMFT-TR-01-105, ALCOM-FT, Patras, May 2001. Appearance: In *Proceedings of the 14th Symposium on Distributed Computing – DISC 2000*, 2001.

[106] Han Hoogeveen and Gerhard J. Woeginger. Some comments on sequencing with controllable processing times. Technical Report ALCOMFT-TR-01-106, ALCOM-FT, Utrecht, May 2001.

[107] Dimitris Fotakis, Sotiris Nikoletseas, Vicky Papadopoulou, and Paul Spirakis. Radiocolorings in Periodic Planar Graphs: PSPACE-Completeness and Efficient Approximations for the Optimal Range of Frequencies. Technical Report ALCOMFT-TR-01-107, ALCOM-FT, Patras, May 2001.

[108] Dimitris Fotakis, Sotiris Nikoletseas, Vicky Papadopoulou, and Paul Spirakis. Hardness Results and Efficient Approximations for Frequency Assignment Problems: Radio Labelling and Radio Coloring. Technical Report ALCOMFT-TR-01-108, ALCOM-FT, Patras, May 2001. Appearance: *Journal of Computers and Artificial Intelligence – CAI 2001*, 2001. To appear.

[109] Dimitris Fotakis, Sotiris Nikoletseas, Vicky Papadopoulou, and Paul Spirakis. NP-completeness Results and Efficient Approximations for Radiocoloring in Planar Graphs. Technical Report ALCOMFT-TR-01-109, ALCOM-FT, Patras, May 2001. Appearance: In *Proceedings of the 25th Symposium Mathematical Foundations of Computer Science – MFCS 2000*. EATCS, 2000.

[110] Frédéric Chyzak. Algorithms Seminar, 1999-2000. Technical Report ALCOMFT-TR-01-110, ALCOM-FT, INRIA, May 2001.

[111] Hans L. Bodlaender. A generic NP-hardness proof for a variant of Graph Coloring. Technical Report ALCOMFT-TR-01-111, ALCOM-FT, Utrecht, May 2001.

[112] Jochen Alber, Hans L. Bodlaender, Henning Fernau, and Rolf Niedermeier. Fixed Parameter Algorithms for PLANAR DOMINATING SET and Related Problems. Technical Report ALCOMFT-TR-01-112, ALCOM-FT, Utrecht, May 2001. APPEARANCE: In Magnús M. Halldórsson, editor, *Proceedings 7th Scandinavian Workshop on Algorithm Theory, SWAT'2000, Lecture Notes in Computer Science 1851*, pages 97–110. Springer-Verlag, 2000.

[113] Hans L. Bodlaender. Necessary edges in k-chordalizations of graphs. Technical Report ALCOMFT-TR-01-113, ALCOM-FT, Utrecht, May 2001.

[114] Dimitrios M. Thilikos, Maria J. Serna, and Hans L. Bodlaender. A constructive linear time algorithm for small cutwidth. Technical Report ALCOMFT-TR-01-114, ALCOM-FT, Barcelona and Utrecht, May 2001. APPEARANCE: In D.T. Lee and Shang-Hua Teng, editors, *Proc. 11th International Symposium on Algorithms And Computation ISAAC '00, Lecture Notes in Computer Science 1969*, pages 192–203. Springer-Verlag, 2000.

[115] Dimitrios M. Thilikos, Maria J. Serna, and Hans L. Bodlaender. A polynomial time algorithm for the cutwidth of bounded degree graphs with small treewidth. Technical Report ALCOMFT-TR-01-115, ALCOM-FT, Barcelona and Utrecht, May 2001.

[116] Burkhard Monien and Robert Preis. Upper Bounds on the Bisection Width of 3- and 4-regular Graphs. Technical Report ALCOMFT-TR-01-116, ALCOM-FT, Paderborn, May 2001. APPEARANCE: In *Proceedings of 26th International Symposium on Mathematical Foundations of Computer Science (MFCS)*, 2001. To appear.

[117] Ioannis Chatzigiannakis, Sotiris Nikoletseas, and Paul Spirakis. An Efficient Communication Strategy for Ad-hoc Mobile Networks. Technical Report ALCOMFT-TR-01-117, ALCOM-FT, Patras, May 2001. APPEARANCE: In *Proceedings of the 20th Symposium on Principles of Distributed Computing – PODC 2001.* ACM, 2001. To appear as Brief Announcement.

[118] J. Diaz, D. Koukopoulos, S. Nikoletseas, M. Serna, P. Spirakis, and D. Thilikos. Stability and non-stability of the FIFO Protocol. Technical Report ALCOMFT-TR-01-118, ALCOM-FT, Barcelona and Patras, May 2001. APPEARANCE: In *Proceedings of the 13th Symposium on Parallel Algorithms and Architectures – SPAA 2001.* ACM, 2001.

[119] Ioannis Chatzigiannakis, Sotiris Nikoletseas, Nearchos Paspallis, Paul Spirakis, and Christos Zaroliagis. An Experimental Study of Basic Communication Protocols in Ad-Hoc Mobile Networks. Technical Report ALCOMFT-TR-01-119, ALCOM-FT, Patras, May 2001. APPEARANCE: In *Algorithm Engineering – WAE 2001.* Lect. Notes in Comp. Science, Springer-Verlag, 2001. To appear.

[120] Marios Mavronicolas and Paul Spirakis. The Price of Selfish Routing. Technical Report ALCOMFT-TR-01-120, ALCOM-FT, Patras, May 2001. APPEARANCE: In *Proceedings of the 33rd Symposium on Theory of Computing – STOC 2001.* ACM, 2001.

[121] Hans L. Bodlaender, Michael J. Dinneen, and Bakhadyr Khoussainov. Relaxed Update and Partition Network Games. Technical Report ALCOMFT-TR-01-121, ALCOM-FT, Utrecht, May 2001.

[122] Sixiang Hou and Han Hoogeveen. The three-machine proportionate flow shop problem with unequal machine speeds. Technical Report ALCOMFT-TR-01-122, ALCOM-FT, Utrecht, May 2001.

[123] Christoph Buchheim and Michael Jünger. Detecting Symmetries by Branch & Cut. Technical Report ALCOMFT-TR-01-123, ALCOM-FT, Cologne, May 2001.

[124] Rasmus Pagh and Flemming Friche Rodler. Cuckoo Hashing. Technical Report ALCOMFT-TR-01-124, ALCOM-FT, Aarhus, May 2001. APPEARANCE: In *Proceedings of ESA 2001*, Lecture Notes in Computer Science. Springer-Verlag, 2001.

[125] Marjan van den Akker, Han Hoogeveen, and Steef van de Velde. Combining column generation and Lagrangean relaxation to solve a single-machine common due date problem. Technical Report ALCOMFT-TR-01-125, ALCOM-FT, Utrecht, May 2001.

[126] Rasmus Pagh and Flemming Friche Rodler. Lossy Dictionaries. Technical Report ALCOMFT-TR-01-126, ALCOM-FT, Aarhus, May 2001. Appearance: In *Proceedings of ESA 2001*, Lecture Notes in Computer Science. Springer-Verlag, 2001.

[127] Piotr Krysta and Anil Kumar. Approximation Algorithms for Minimum Size 2-Connectivity Problems. Technical Report ALCOMFT-TR-01-127, ALCOM-FT, MPI, May 2001. Appearance: In *Proceedings of The 18th International Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 2010 of *Lecture Notes in Computer Science*, pages 431–442. Springer Verlag, Berlin, 2001.

[128] Rolf Fagerberg, Rune E. Jensen, and Kim S. Larsen. Search Trees with Relaxed Balance and Near-Optimal Height. Technical Report ALCOMFT-TR-01-128, ALCOM-FT, Aarhus, May 2001. Appearance: In *Proceedings of WADS 2001*, Lecture Notes in Computer Science. Springer-Verlag, 2001.

[129] Hans L. Bodlaender, Richard B. Tan, and Jan van Leeuwen. Finding a Δ-regular Supergraph of Minimum Order. Technical Report ALCOMFT-TR-01-129, ALCOM-FT, Utrecht, May 2001.

[130] Gerth Stølting Brodal, Rolf Fagerberg, Christian N. S. Pedersen, and Anna Östlin. The Complexity of Constructing Evolutionary Trees Using Experiments. Technical Report ALCOMFT-TR-01-130, ALCOM-FT, Aarhus, May 2001. Appearance: In *Proceedings of ICALP 2001*, Lecture Notes in Computer Science. Springer-Verlag, 2001.

[131] Gerth Stølting Brodal, Rolf Fagerberg, and Christian N. S. Pedersen. Computing the Quartet Distance Between Evolutionary Trees in Time $O(n \log^2 n)$. Technical Report ALCOMFT-TR-01-131, ALCOM-FT, Aarhus, May 2001.

[132] Artur Czumaj and Christian Sohler. Soft Kinetic Data Structures. Technical Report ALCOMFT-TR-01-132, ALCOM-FT, Paderborn, May 2001. Appearance: In *Proceedings of the 12th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 865–872, 2001.

[133] Artur Czumaj and Christian Sohler. Testing Hypergraph Coloring. Technical Report ALCOMFT-TR-01-133, ALCOM-FT, Paderborn, May 2001. Appearance: In *Proceedings of the 28th International Colloquium on Automata, Languages and Programming (ICALP)*, 2001. To appear.

[134] Michael Wand, Matthias Fischer, Ingmar Peter, Friedhelm Meyer auf der Heide, and Wolfgang Straßer. The Randomized z-Buffer Algorithm: Interactive Rendering of Highly Complex Scenes. Technical Report ALCOMFT-TR-01-134, ALCOM-FT, Paderborn, May 2001. Appearance: In *Proceedings of the ACM Computer Graphics (SIGGRAPH)*, 2001. To appear.

[135] J. Díaz, M. Serna, and D.M. Thilikós. Counting List H-Colorings and Variants. Technical Report ALCOMFT-TR-01-135, ALCOM-FT, Barcelona, June 2001.

[136] J. Díaz, J. Petit, and M. Serna. A survey on graph layout problems. Technical Report ALCOMFT-TR-01-136, ALCOM-FT, Barcelona, June 2001.

[137] C. Àlvarez, R. Cases, J. Díaz, J. Petit, and M. Serna. Communication tree problems. Technical Report ALCOMFT-TR-01-137, ALCOM-FT, Barcelona, June 2001. Appearance: In J. Rolim, editor, *ICALP Workshops 2000*, volume 8 of *Proceedings in Informatics*, pages 99–110, Canada, 2000. Carleton Scientific. This is a preliminary version of the full paper.

[138] J. Petit. Hamiltonian cycles in faulty random geometric networks. Technical Report ALCOMFT-TR-01-138, ALCOM-FT, Barcelona, June 2001.

[139] A. Stewart, M. Clint, J. Gabarró, and M. Serna. Towards refining BSP programs into loosely coupled distributed systems. Technical Report ALCOMFT-TR-01-139, ALCOM-FT, Barcelona, June 2001. Appearance: In *Proceedings of The EURO-PAR 2001 conference*, Lecture Notes in Computer Science, 2001.

[140] J. Díaz, J. Petit, and M. Serna. Faulty random geometric networks. Technical Report ALCOMFT-TR-01-140, ALCOM-FT, Barcelona, June 2001. Appearance: *Parallel Processing Letters*, 10(4):343–357, 2001.

[141] Rasmus Pagh and Jakob Pagter. Optimal Time-Space Trade-Offs for Non-Comparison-Based Sorting. Technical Report ALCOMFT-TR-01-141, ALCOM-FT, Aarhus, June 2001.

[142] J. Díaz, M. Serna, and D. M. Thilikos. Counting H-colorings of Partial k-trees. Technical Report ALCOMFT-TR-01-142, ALCOM-FT, Barcelona, June 2001. Appearance: In *Proceedings of The 7th Annual International Computing and Combinatorics Conference, COCOON 2001*, Lecture Notes in Computer Science, 2001.

[143] J. Díaz, M. Serna, and D. M. Thilikos. The Complexity of Parameterized H-Colorings: A survey. Technical Report ALCOMFT-TR-01-143, ALCOM-FT, Barcelona, June 2001.

[144] Rune Bang Lyngsø and Christian Nørgaard Storm Pedersen. Complexity of Comparing Hidden Markov Models. Technical Report ALCOMFT-TR-01-144, ALCOM-FT, Aarhus, June 2001.

[145] Arun Jagota, Rune Bang Lyngsø, and Christian Nørgaard Storm Pedersen. Comparing Stochastic Models. Technical Report ALCOMFT-TR-01-145, ALCOM-FT, Aarhus, June 2001.

[146] Graham Cormode, S. Muthukrishnan, and Süleyman Cenk Ṣahinalp. Permutation Editing and Matching via Embeddings. Technical Report ALCOMFT-TR-01-146, ALCOM-FT, Warwick, June 2001. Appearance: In F. Orejas, P.G. Spirakis, and J. van Leeuwen, editors, *Proceedings of 28th International Colloquim on Automata, Languages and Programming*, volume 2076 of *Lecture Notes in Computer Science*, July 2001.

[147] C. Demetrescu. Fully Dynamic Algorithms for Path Problems on Directed Graphs. Technical Report ALCOMFT-TR-01-147, ALCOM-FT, Rome, June 2001. Appearance: PhD thesis, Department of Computer and Systems Science, University of Rome "La Sapienza", report XIII-01-1, February 2001.

[148] C. Demetrescu and I. Finocchi. Removing Cycles for Minimizing Crossings. Technical Report ALCOMFT-TR-01-148, ALCOM-FT, Rome, June 2001. Appearance: In B.E. Moret and A.V. Goldberg, editors, *Proc. of the 2nd International Workshop on Algorithm Engineering and Experiments (ALENEX'00)*, pages 171–182, 2000.

[149] C. Demetrescu, I. Finocchi, G.F. Italiano, and S. Naeher. Visualization in Algorithm Engineering: Tools and Techniques. Technical Report ALCOMFT-TR-01-149, ALCOM-FT, Rome, June 2001.

[150] C. Demetrescu and G.F. Italiano. Fully Dynamic All Pairs Shortest Paths with Real Edge Weights. Technical Report ALCOMFT-TR-01-150, ALCOM-FT, Rome, June 2001. Appearance: In *Proc. of the 42nd IEEE Annual Symposium on Foundations of Computer Science (FOCS'01), Las Vegas, Nevada*, 2001.

[151] M.J. Blesa and F. Xhafa. A Skeleton for the Tabu Search Metaheuristic with Applications to Problems in Software Engineering. Technical Report ALCOMFT-TR-01-151, ALCOM-FT, Barcelona, June 2001. APPEARANCE: In *SEMINAL: Software Engineering using Metaheuristic INnovative ALgorithms*, pages 23–31, 2001. Complement of the *Proceedings of the 23rd International Conference on Software Engineering*.

[152] M.J. Blesa, Ll. Hernàndez, and F. Xhafa. Parallel Skeletons for Tabu Search Method. Technical Report ALCOMFT-TR-01-152, ALCOM-FT, Barcelona, June 2001. APPEARANCE: In *8th International Conference on Parallel and Distributed Systems*. IEEE Computer Society Press, 2001. Shorter version to appear.

[153] Funda Ergun, S. Cenk Sahinalp, Jonathan Sharp, and Rakesh K. Sinha. Biased Dictionaries with Fast Insert/Deletes. Technical Report ALCOMFT-TR-01-153, ALCOM-FT, Warwick, June 2001. APPEARANCE: In *Proceedings of STOC*, 2001.

[154] Funda Ergun, S. Cenk Sahinalp, Jonathan Sharp, and Rakesh K. Sinha. Biased Skip Lists for Highly Skewed Access Patterns. Technical Report ALCOMFT-TR-01-154, ALCOM-FT, Warwick, June 2001. APPEARANCE: In *Proceedings of ALENEX*, 2001.

[155] M.J. Blesa, Ll. Hernandez, and F. Xhafa. Parallel Skeletons for Tabu Search Method Based on Search Strategies and Neighborhood Partition. Technical Report ALCOMFT-TR-01-155, ALCOM-FT, Barcelona, June 2001.

[156] A. Stewart, M. Clint, and J. Gabarró. Algebraic rules for reasoning about BSP programs. Technical Report ALCOMFT-TR-01-156, ALCOM-FT, Barcelona, June 2001. APPEARANCE: In *CMPP 2000, Second International Workshop on Constructive Methods for Parallel Programming*, pages 35–51, 2000.

[157] Jordi Petit. Layout Problems. Technical Report ALCOMFT-TR-01-157, ALCOM-FT, Barcelona, June 2001.

[158] Mireille Régnier, Alexandre Lifanov, and Vsevolod Makeev. Three variations on word counting. Technical Report ALCOMFT-TR-01-158, ALCOM-FT, INRIA, June 2001.

[159] Alain Denise, Mireille Régnier, and Mathias Vandenbogaert. Assessing statistical significance of overrepresented oligonucleotides. Technical Report ALCOMFT-TR-01-159, ALCOM-FT, INRIA, June 2001.

[160] Jan van Leeuwen and Jiri Wiedermann. The Turing Machine Paradigm in Contemporary Computing. Technical Report ALCOMFT-TR-01-160, ALCOM-FT, Utrecht, June 2001. APPEARANCE: In B. Enquist and W. Schmidt, editors, *Mathematics Unlimited - 2001 and Beyond*, pages 1139–1155. Springer-Verlag, 2001.

[161] Jan van Leeuwen and Jiri Wiedermann. A Computational Model of Interaction in Embedded Systems. Technical Report ALCOMFT-TR-01-161, ALCOM-FT, Utrecht, June 2001.

[162] Jan van Leeuwen and Jiri Wiedermann. On Algorithms and Interaction. Technical Report ALCOMFT-TR-01-162, ALCOM-FT, Utrecht, June 2001. APPEARANCE: In *Mathematical Foundations of Computer Science 2000, 25th Int. Symposium*, volume 1893 of *Lecture Notes in Computer Science*, pages 99–113, 2000. Revised version: 'A computational model of interaction in embedded systems', Techn. Rep. UU-CS-2001-02, Institute of Information and Computing Sciences, Utrecht University, 2001.

[163] Jan van Leeuwen and Jiri Wiedermann. On the Power of Interactive Computing. Technical Report ALCOMFT-TR-01-163, ALCOM-FT, Utrecht, June 2001. APPEARANCE: In *Theoretical Computer Science, Exploring New Frontiers of Theoretical Informatics (IFIP TCS 2000)*, volume 1872 of *LNCS*, pages 619–623, 2000. Revised and extended version: 'The Turing machine paradigm in contemporary computing', in: B. Enquist and W. Schmidt (Eds), Mathematics Unlimited - 2001 and Beyond, Springer-Verlag, 2001, pp 1139-1155.

[164] Lene Monrad Favrholdt and Morten Nyhave Nielsen. On-Line Edge-Coloring with a Fixed Number of Colors. Technical Report ALCOMFT-TR-01-164, ALCOM-FT, Aarhus, June 2001. APPEARANCE: *FSTTCS: Foundations of Software Technology and Theoretical Computer Science*, 20, 2000.

[165] M.J. Blesa, P. Moscato, and F. Xhafa. A Memetic Algorithm for the Minimum Weighted k-Cardinality Tree Subgraph Problem". Technical Report ALCOMFT-TR-01-165, ALCOM-FT, Barcelona, June 2001. APPEARANCE: In *4th Metaheuristics International Conference*. Kluwer, 2001. To appear.

[166] Fedor V. Fomin and Dimitrios M. Thilikos. On the Monotonicity of Games Generated by Symmetric Submodular Functions. Technical Report ALCOMFT-TR-01-166, ALCOM-FT, Barcelona, June 2001. APPEARANCE: *Discrete Applied Mathematics*. To appear.

[167] Salvador Roura. A New Method for Balancing Binary Search Trees. Technical Report ALCOMFT-TR-01-167, ALCOM-FT, Barcelona, June 2001. APPEARANCE: In *International Colloquium on Automata, Languages and Programming (ICALP)*, 2001.

[168] Marni Mishna. Attribute Grammars and Automatic Complexity Analysis. Technical Report ALCOMFT-TR-01-168, ALCOM-FT, INRIA, June 2001.

[169] Friedhelm Meyer auf der Heide and Rolf Wanka. Parallel Bridging Models and Their Impact on Algorithm Design. Technical Report ALCOMFT-TR-01-169, ALCOM-FT, Paderborn, June 2001. APPEARANCE: In *Proceedings of the International Conference on Computational Science (ICCS), Vol. II*, pages 628–637, 2001.

[170] Ricard Gavaldà and Denis Thérien. Learning expressions over monoids. Technical Report ALCOMFT-TR-01-170, ALCOM-FT, Barcelona, June 2001. APPEARANCE: In *Proc. 18th Intl. Symposium on Theoretical Aspects of Computer Science (STACS'01)*, volume 2010, pages 283–293. Springer-Verlag LNCS, 2001.

[171] Lars Arge and Jakob Pagter. I/O-Space Trade-Offs. Technical Report ALCOMFT-TR-01-171, ALCOM-FT, Aarhus, June 2001. APPEARANCE: In *Proceedings of the 7th Scandinavian Workshop on Algorithm Theory (SWAT 2000)*, volume 1851 of *Lecture Notes in Computer Science*, pages 448–461. springer, Berlin, 2000.

[172] Jakob Pagter. On Ajtai's Lower Bound Technique for R-way Branching Programs and the Hamming Distance Problem. Technical Report ALCOMFT-TR-01-172, ALCOM-FT, Aarhus, June 2001.

[173] Peter Sanders and Jop Sibeyn. A Bandwidth Latency Tradeoff for Broadcast and Reduction. Technical Report ALCOMFT-TR-01-173, ALCOM-FT, MPI, August 2001. APPEARANCE: In *6th Euro-Par*, number 1900 in LNCS, pages 918–926, 2000.

[174] Berthold Vöcking. Almost optimal permutation routing on hypercubes. Technical Report ALCOMFT-TR-01-174, ALCOM-FT, MPI, August 2001. APPEARANCE: In *ACM Symposium on Theory of Computing*, pages 530–539, 2001.