Second Annual Progress Report

ALCOM-FT

Algorithms and Complexity
Future Technologies

Project No. IST-1999-14186

June 2002

Summary

This is the second annual progress report for the ALCOM-FT project, supported by the European Commission as project number IST-1999-14186 under the Future and Emerging Technologies part of the IST programme of the Fifth Framework. The report covers the period June 1, 2001 to May 31, 2002.

ALCOM-FT brings together eleven 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 eleven 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
11	Department of Computer Science, University of Cyprus	Cyprus

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

During the second year of the project, the work has been carried out basically as planned. Proper coordination of the project has been ensured by the efforts of the Consortium Board, the Work Package Leaders, and the Coordinator. Dissemination of the work done has been ensured by the publication of scientific reports, by summer schools, by industrial talks, by an algorithmic competition, and by the maintenance of websites for the entire project as well as for several deliverables within the project.

The second year of the project also saw the addition of one further site to the ALCOM-FT consortium, namely the algorithmics group at the University of Cyprus, headed by Marios Mavronicolas. The addition of the group will strengthen the project within Work Packages 2 and 4.

The eleven deliverables relating to the second year of the project are listed in Table 2. The largest deliverable is D1, which represents a total of 158 scientific reports, produced during the second project year. All reports are available online. Deliverables D13, D16, D17, D18, D20, and D22 have been delivered as planned whereas deliverables D14, D15, D19 and D21 are delayed, mainly due to difficulties with filling the relevant programming positions. From last year, deliverable D10 still remains to be completed. It is a volume in the *Lecture Notes in Computer Science* series of *Springer Verlag*, and has been delayed due to the time involved in the editorial process. It will appear in the latter half of 2002.

The rest of this document is organized into four parts. The first describes the progress achieved

within the various work packages of the project. The second lists significant events during the year for each participating site. The third tabulates the resources used in each of the tasks of the project. The fourth lists the scientific reports produced during the second year of the project.

No.	*Deliverable*	*Month*
D1	Research reports	Cont.
D13	Testbed for experimental algorithmics (internal release)	18
D14	PR web-pages for algorithmics	18
D15	External memory experimental platform (prototype)	24
D16	Data mining demonstration package	24
D17	Distributed algorithmic engineering software package (beta version)	24
D18	BSP-style library for dynamic distributed environments (beta version)	24
D19	Production and transportation planning software prototype	24
D20	Description language for high level constraints in optimization	24
D21	Testbed for experimental algorithmics (external release)	24
D22	Second progress report	24

Table 2: The deliverables relating to the second year.

Contents

Part I Work Packages

WP 1: Massive Data Sets

Progress in this work package is documented in 18 technical reports and by the preliminary implementation of some of the new algorithms contributed by the project. The relevant reports are [22–24, 43, 46, 48, 53, 57, 84, 90, 108, 117, 120, 136, 138, 148–150, 154].

About half of them correspond to the core tasks of this work package, the others being either side developments that were felt necessary for the progress of the theory, or applications to other domains, with priority to the problems tackled in the other work packages of the project. We describe in more detail those within the core tasks of the work package: algorithms for external memory (which we for clarity have structured into a section on primitives for the handling of parallel external memory and a section on graph algorithms), algorithms for data mining, and algorithms that are designed to work on cache memories but in a form that is optimal independently of the cache size (so-called cache-oblivious algorithms). An additional last section summarizes briefly the other contributions of the project.

Parallel external memory

Massive data sets need not only to be processed efficiently but also to be stored in a flexible manner. Such a storage scheme has to permit fast access to data elements even if the access pattern is not known in advance. This is necessary because even the large efforts to find good external memory algorithms for fundamental problems (like graph or geometric problems) did not result in many practical solutions for real-life problems.

Current storage systems lack the flexibility to adapt to a growing space demand. One way of meeting capacity and performance requirements is the usage of dedicated storage networks. These networks consist of a number of storage devices (e.g. disk drives) which are connected by an arbitrary network. The arising problem is the efficient mapping of data elements to storage devices. Very simple storage schemes are sufficient to solve the homogeneous (i.e. all storage devices are the same) static case but are insufficient for heterogeneous (i.e. storage devices have different capacity and/or transfer rate) and adaptive (i.e. storage devices may enter or leave the system) settings.

We are looking for solutions to the general problem needing to perform an optimal number of replacements in the adaptive case. In the previous year we proposed a solution to the homogeneous problem. This scheme worked in two phases. Firstly, the full address space (all possible data elements) was mapped to the real $[0, 1)$ interval using a pseudo-random hash function. In the second phase, this interval is cut into sub-intervals which are mapped (along with all data elements falling into that interval) to storage devices via a so-called assimilation function. This function ensures that only the optimal number of data elements is redistributed when the number of storage devices changes in the expected case. Furthermore, this function defines the access to data elements. One has simply to follow the number of redistributions done when the number of storage devices changes from 1 to n. We have shown that only $O(\log n)$ simple computations are sufficient to access a data element.

Unfortunately, this scheme is not easily extendible to the heterogeneous case. In the previous year we analysed the natural extension in which we define new levels of unsaturated storage devices as soon as one storage device gets saturated. This increases the data access by a factor of $O(\log m)$, i.e. the logarithm of the number of data elements. Recently, we developed two new placement schemes: SHARE and SIEVE [43]. Both of them are the first storage algorithms that permit an optimal adaptation to changing space requirements. The SHARE algorithm extends the consistent hashing approach so that heterogeneous requirements can be handled. Instead of casting points representing the storage devices into a $[0, 1)$ interval, range-intervals of length according the requirements are used. SIEVE has a number of levels in each of which a number of

data elements is mapped to storage devices in a sieve-like fashion. In each level the data elements are hashed using a pseudo-random hash function.

Graph algorithms

We continued to consider the classical graph traversal problems breadth-first search (BFS), depth-first search (DFS) and single-source shortest paths (SSSP) for large inputs. Using parallel computing we solve the problems faster and/or take advantage of the aggregated internal memory. On the other hand, for external-memory graph traversal, we show novel ways to avoid unstructured disk accesses.

In [138] we provide a new bucket-based parallel SSSP algorithm. We study its average-case complexity assuming arbitrary directed graphs with n nodes, m edges, and independent random edge weights uniformly distributed in $[0, 1]$; it achieves average-case running time $T = O(\log^2 n \cdot \min_i\{2^i \cdot \mathbf{E}[\mathcal{L}] + |V_i|\})$ using $O(n + m + T)$ work on a PRAM where $\mathcal{L}$ denotes the maximum shortest-path weight and $|V_i|$ is the number of graph vertices with in-degree at least 2^i. All previous algorithms either required more time or more work. The minimum performance gain is a logarithmic factor improvement; on certain graph classes, accelerations by factors of more than $n^{0.4}$ can be achieved. The algorithm also allows adaptation to distributed memory machines.

The best previously known external-memory algorithms for BFS and DFS required $\Omega(n)$ I/Os on general graphs. This is even true for *semi-external computing*, where it is assumed that node-related data can be stored internally. In [154] we give a set of heuristics which allow to perform semi-external DFS for directed graphs in practice. Our approach uses blocked I/O in order to successively transform a spanning tree of the graph into a valid DFS tree; the tree itself is always kept in internal memory. The heuristics have been applied to graphs with very different graph properties, including models of the web graph and some call graphs from ATT. The program is typically about 100 times faster than the best alternative, a factor that will further increase with future technology developments.

In [136] we present the first external-memory BFS algorithm that achieves $o(n)$ I/Os on arbitrary undirected graphs. On a machine with main-memory of size M, D parallel disks, and block size B, our new algorithm requires only $O\left(\sqrt{\frac{1}{D \cdot B}} + \sqrt{\frac{n}{m}}) \cdot \frac{n+m}{\sqrt{D \cdot B}} \cdot \log_{M/B} \frac{n+m}{B}\right)$ I/Os. Hence, for $m = O(n)$, they improve upon the I/O-performance of the best previous algorithm by nearly a factor of $\sqrt{D \cdot B}$. Our approach is fairly simple and we conjecture it to be practical. We also give an improved single-source shortest-paths algorithm for small integer edge weights.

Sampling approaches to data mining

A central data mining problem is the computation of frequent item sets from a transactional dataset. In combinatorial terms, the dataset is actually a hypergraph, and computing frequent sets is closely related to the problem of computing minimal transversals. In the previous year we designed a new frequent sets algorithm. Besides being efficient, it has the main advantage of combining naturally with adaptive sampling schemes. With respect to the widely used batch sampling, adaptive sampling maintains the same theoretically well-founded statistical guarantee of representativeness; simultaneously, in the vast majority of practical cases, it requires a reasonably small sample size. This is not the case with batch sampling, which gives extremely demanding sample sizes. The standard schemes for frequent sets algorithms are unable to take profit from the advantage of adaptive sampling. Previous work along these lines has continued its development into a unified framework [108] to study adaptive sampling schemes, ready for implementation and allowing also to prove limitations of the approach in the form of lower bounds. Two natural ways of applying adaptive sampling on our frequent sets algorithms have been studied in [22], where also a third one has been designed that performs better than both in a broad set of practical experiments.

With the additional ingredient of letting the underlying probability distribution, according to which the sampling is done, evolve in time, we have been able to apply sampling strategies

to the training of Support Vector algorithms for classification problems; these algorithms are becoming another central topic in the datamining community. Algorithms designed in the first year (under implementation) have appeared in the meantime in major conferences of Learning Theory and of Data Mining. Additionally, during the second year we have been able to extend the randomized sampling-based scheme to the regression problem [24]. Work remains to be done on tackling problems of high dimensionality, where our randomized algorithmic techniques may prove effective as well.

In most cases, the next application of the output of a frequent sets algorithm is to obtain the association rules from it. Current formulations of the association rules problem causes two problems: on the one hand, not all interesting facts are representable in this way, but, on the other, the number of association rules obtained, even from small datasets, is already too large to be of use, since its usefulness depends on human inspection. We proposed in the past a scheme to improve the expressibility of association rules by including negations, but with no algorithmic efficiency loss. This year, we have studied methods to reduce the number of obtained rules to manageable sizes with either no information loss at all, or by comparison with a standard statistical correlation factor (the Pearson coefficient). The first approach has been successfully applied to association rules with negations [23], and the second has focused on standard association rules but has proposed to substitute a multi-step pruning method for the single-step methods currently in use, with highly improved results [57].

Cache oblivious algorithms

Recently, the concept of *cache oblivious* algorithms has been introduced by Frigo et al. In essence, this designates algorithms optimized in the I/O model, except that one optimizes to a block size B and a memory size M which are *unknown*. This seemingly simple change has significant consequences: since the analysis holds for any block and memory size, it holds for *all* levels of the memory hierarchy. In other words, by optimizing an algorithm to one unknown level of the memory hierarchy, it is optimized to each level automatically. Also, the characteristics of the memory hierarchy do not need to be known to the algorithm, thereby increasing portability, which is important in situations such as the production of software libraries and code delivered over the web.

pointer to this at the end.

Sorting is a fundamental task in information processing, and is also one of the few currently known tools for designing cache oblivious algorithms. Improving the practicality of sorting algorithms is therefore important. In [46], we do this by simplifying the cache oblivious sorting algorithm Funnelsort of Frigo et al., and showing how to relax the "tall cache assumption" from $M \geq B^2$ to $M \geq B^{1+\epsilon}$. Based on this simplified algorithm, we then give cache oblivious algorithms for a number of problems in computational geometry.

Arge et al. recently presented the first optimal cache oblivious priority queue, and demonstrated the importance of this result by providing the first cache oblivious algorithms for graph problems. Their structure uses cache oblivious sorting and selection as subroutines. In [47], we give an alternative optimal cache oblivious priority queue based only on binary merging. We also show that our structure can be made adaptive to different usage profiles.

Work on cache oblivious search trees [48] is described in WP5.

Other work

We continued work on approximate string matching in text databases using filtering methods based on gapped q-grams. Efficient use of these filters requires optimization of a parameter called the threshold. A generic version of this optimization problem together with algorithms for solving it by reduction to a constrained shortest path problem are described in [120]. A class of filters for the widely used Levenshtein string distance model is introduced and studied in [53].

In [117] and [90] we discuss space-efficient dictionary data-structures with applications in external memory. Hard-disk and cache performance are analyzed under specific access patterns and

data distributions in [148] and [150].

When processing large volumes of data, we also naturally get the problem of exchanging large amounts of data between several processes. The paper [149] explains how to implement one of the most common communication patterns on hierarchically interconnected systems.

Finally, [84] presents an average analysis of the complexity of multidimensional search algorithms in very large data structures.

WP 2: Networks and Communication

The principal aim of the work within WP2 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 second year of the project addressed the following topics:

- Modelling of telecommunication networks,
- Algorithms for wireless networks,
- Adversarial queueing theory,
- Selfish routing and game–theoretic models,
- Load balancing,
- A platform for efficient computing in distributed dynamic environments.

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

Modelling of telecommunication networks

In this line of research, INRIA 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.

The problems investigated during the second year of the project are the following:

a) A complete investigation of the behavior of Additive Increase Multiplicative Decrease Algorithms used in congestion avoidance mechanisms of communication protocols like TCP [86, 113, 114],

b) Admission control algorithms for Diffserv networks [87, 118],

c) Further studies concerning limit theorems of rescaled stochastic processes describing communication networks without buffering [106].

TCP Protocol. The congestion avoidance phase of TCP (Transmission Control Protocol) networks is an example of an AIMD algorithm. An AIMD algorithm can be quickly described as follows: a source transmits W packets in the network and if an acknowledgement of reception is received by the source for all the packets, it then sends $W + 1$ packets (Additive Increase). Otherwise, a packet loss is detected. The source then sends $\lfloor \delta W \rfloor$ packets (Multiplicative Decrease). The variable W is called *congestion window size*. The analysis is focused on the evolution of the variable W for a very long connection: a source is sending a large file to a destination and the network loses packets from time to time. Two different models have been considered for the loss process.

Model 1. Packet losses are independent [86].

Model 2. Packet losses are correlated, losses occur in groups [113, 114].

Model 1 is, of course, not completely realistic since packet losses are due to overflows of buffers. If a packet is lost, it is very likely that some of the consecutive packets will also be lost so that independence cannot be strict. It turns out, however, as shown in Model 2, that this case provides useful lower bounds for the performance of the algorithm. Moreover, the convergence results obtained in this case can be used directly in Model 2.

In both cases it is shown that the congestion window size at equilibrium is equivalent to $\overline{W}/\sqrt{\alpha}$, when the loss rate α tends to 0, where $\overline{W}$ is a random variable. This is the first study of TCP where rigorous convergence results are obtained.

Model 1. Independent case. This case is handled by renormalizing the time by a factor $1/\sqrt{\alpha}$ and the variable W by $\sqrt{\alpha}$. It is very helpful to consider also time in the limit theorems (instead of only the equilibrium). A limiting Markov process is obtained in this manner, its analysis gives most of the results concerning the behavior of the algorithm (in particular, the density of $\overline{W}$ and the throughput of the connection).

Model 2. In this case the distribution of the variable $\overline{W}$ is much more intricate than for the independent case. On the analytical side, q-calculus is the natural framework to study the density of this variable. In some cases this density is related to q-hypergeometric functions. On the probabilistic side, a very interesting relation with stochastic models used in mathematical finance and statistical physics has been shown. Random variables called exponential functionals play a central role in our analysis. These variables are the analogue of the variables used in mathematical finance in a Brownian framework.

By using concave ordering of random variables, it has also been proved that, for a fixed loss rate, correlation of packet losses improves the performance of the protocol. The higher the variability of the loss process is, the better the throughput. In particular, the independent loss model underestimates the real performance of TCP.

Another line of research related to the TCP protocol was studied by Aarhus. In point-to-point communication, small messages, such as TCP acknowledgments, may be bundled and sent as one packet. Whereas waiting for additional messages causes delay for the ones already waiting, sending immediately can delay later messages, since networks require a certain distance between packets. In [105], deterministic and randomized algorithms are characterized, using a model of communication which is more realistic than the ones previously used for this question.

Connection admission control At the entrance of the network such an algorithm determines whether a request for connection can be accepted or not [87]. Connections are classified in several types according to their requirements (throughput, load, etc.). The constraints are the following:

- Accept the maximum number of connections (i.e. the number of subscribers for Assured Forwarding class of Traffic),
- Guarantee that any accepted connection will have a probability of a packet loss less than ε. (Guaranteed quality of service.)

When there is only Assured Forwarding traffic, a number called *effective bandwidth* α_i is associated to each type of connection $i \in I$. If the sum of all the effective bandwidths of the *accepted* connections is S, a new request for connection of type i_0 is accepted only if $\alpha_{i_0} + S$ is less than some constant H determined once and for all. This algorithm is known to be very efficient in this context.

In a Diffserv context however, there is also an additional class of traffic, the Expedite Forwarding Traffic, which has absolute priority over every other traffic. Connection admission control algorithms for Assured Forwarding Traffic have also been designed in this context. Basically, the algorithms considered up to now use also effective bandwidths: they ignore the Expedite Forwarding Traffic but reduce the constant H by ρ times the average load of the Expedite Forwarding

Traffic. If these algorithms are satisfactory in some cases, sometimes they *underestimate* severely the high priority traffic. In some cases, this can lead to a probability of a packet loss much greater than ε, thereby violating the advertised Guaranteed quality of service. (A similar reduced load algorithm is analyzed in [118].)

By considering a new stochastic model as a stochastic upper bound of the initial model, INRIA proposed a new connection admission algorithm in [87]. It uses also effective bandwidths, but it includes the high priority traffic so that quality of service is indeed guaranteed. Moreover, our approach provides explicit bounds for the probability of a packet loss (instead of an equivalent of its logarithm, as it is done usually in the literature).

Loss networks The research of INRIA in this topic [106] is a continuation of the work described in the last year's report ALCOMFT-TR-01-44. They proved that a degenerated functional central limit theorem holds for some saturated loss networks.

Algorithms for wireless networks

In this line of research, we have considered different models for state–of–the–art or future wireless networks and, in each case, we propose new algorithms for resource optimization.

Frequency assignment Patras in [11] studies the Radiocoloring Problem (RCP) on hierarchically specified planar graphs. Hierarchical graphs, besides their inherent combinatorial interest, are motivated by large, structured, wireless communication networks. It is proved that RCP is PSPACE-complete for a wide class of hierarchical planar graphs. Also, efficient approximation algorithms for RCP on such graphs are given.

Ad-hoc mobile networks The work [60] of Patras continues the research on communication protocols for ad-hoc mobile networks. By using combinatorial tools (e.g. random walks) and experiments, a methodology for the average case analysis of protocols for such networks is given, leading to rigorous proofs of correctness and a priori guaranteed performance bounds. Also, [59] focuses on adaptive implementations of such protocols in ad-hoc mobile networks that are highly changing.

Smart Dust networks Smart Dust is a set of a vast number of ultra-small fully autonomous computing and communication devices, with very restricted energy and computing capabilities, that co-operate to quickly and efficiently accomplish a large sensing task. Smart Dust can be very useful in practice e.g. in the local detection of a remote crucial event and the propagation of data reporting its realization. The work [61] of Patras opens the basic algorithmic research on smart dust systems, by providing a simple and realistic model for smart dust, and by giving a set of simple and very efficient protocols for local detection and propagation along with a rigorous average case performance analysis.

Energy optimal routing in radio networks Since stations in mobile networks often run on battery power, low energy consumption is essential for efficient communication. Even one of the simplest questions turns out to be quite challenging here: If we know all positions of the stations and there are no obstacles, what is the optimal route for a message to be sent from a sender to a receiver? Finding an optimal route requires finding a shortest path in an implicitly defined complete graph. Seemingly this takes quadratic time which is prohibitive for large networks. Members of MPI have proposed the first near linear time algorithm (time $O(n^{1+\epsilon})$ for any constant $\epsilon > 0$) by exploiting the geometric structure of the problem [32]. In a further simplified variant of the model, path queries are answered in logarithmic time after $O(n \log n)$ time preprocessing using linear space. The same algorithms can be used as heuristics for reducing the amount of interference between multiple concurrent communications.

Parallel scheduling in next generation wireless networks Next generation 3G/4G wireless data networks allow multiple codes (or channels) to be allocated to a single user, where each code can support multiple data rates. Providing fine-grained QoS to users in such networks poses the two dimensional challenge of assigning both power (rate) and codes to every user. This gives rise to a new class of parallel scheduling problems. Members of Rome in [31] abstract general downlink scheduling problems suitable for proposed next generation wireless data systems. They present provable results on the algorithmic complexity of minimizing the maximum response time scheduling problems and provide simple online algorithms whose performance is analyzed using resource augmented competitive analysis. They also perform an experimental study with realistic data of channel conditions and user requests to show that our algorithms are more accurate than what worst case analysis shows.

Adversarial queueing theory

Adversarial queueing theory is a currently major arena in which to study the *worst-case* performance of routing and contention-resolution protocols for communication networks. The worst-case setting is much more natural and realistic than settings making probabilistic assumptions about the injection of packets into the network and their associated delays (with classical queuing theory being the prime example for the latter settings). Research activity in this arena has attempted to classify protocols and networks according to the *stability* inherited to the network - the property that the number of packets in the network remains *bounded* at certain rates of injection. Cyprus and Patras have recently made substantial progress in several stability-related questions. First, they have shown the smallest (known) injection rate (0.739) that leads to instability when used on a certain FIFO network [124]. Second, they have shown the largest (known) injection rate, as a function of parameters of any particular network, that guarantees stability when used on *any* FIFO network [125]. Third, they have shown that certain contention-resolution protocols that are stable in isolation are not so when composed together in a heterogeneous network with sufficiently large injection rate [124]. Barcelona has provided a complete characterization of universal stability in different path-packet models. For directed networks, this was done in [9], and for undirected networks in [10]. Their result show that testing for universal stability can be done in polynomial time. The results also show that in all the considered cases testing for universal stability is equivalent to testing stability under the NTG-LIS protocol. These results push much forward the current front of research on stability issues in the context of Adversarial Queueing Theory.

Selfish routing and game–theoretic models

In [69] MPI studies the problem of traffic routing in non-cooperative networks. In such networks, users may follow selfish strategies to optimize their own performance measure and therefore their behavior does not have to lead to optimal performance of the entire network. In particular, the worst-case coordination ratio is investigated, which is a game theoretic measure aiming to reflect the price of selfish routing.

Following a line of previous joint work of Patras and Cyprus, members of MPI focus on the most basic networks consisting of parallel links with linear latency functions. Their main result is that the worst-case coordination ratio on m parallel links of possibly different speeds is

$$\Theta\left(\frac{\log m}{\log\log\log m}\right) .$$

In fact, an exact description of the worst-case coordination ratio was given depending on the number of links and the ratio of the speed of the fastest link over the speed of the slowest link. For example, for the special case in which all m parallel links have the same speed, it has been proved that the worst-case coordination ratio is $\Gamma^{(-1)}(m) + \Theta(1)$ with Γ denoting the Gamma (factorial) function. These bounds entirely resolve an open problem posed recently by Koutsoupias and Papadimitriou.

In [67], MPI presents the first thorough study of the price of selfish routing under general, monotone families of cost functions and for cost functions from Queueing Theory. In particular, a precise characterization of cost functions having a bounded/unbounded coordination ratio is given. For example, cost functions that describe the expected delay in queueing systems have an unbounded coordination ratio.

Special focus lies on cost functions describing the behavior of Web servers that can open only a limited number of TCP connections. In particular, the performance of queueing systems that serve all incoming requests is compared to the performance of servers that reject requests in case of overload. It is concluded that queuing systems without rejection cannot give any reasonable guarantee on the expected delay of requests under selfish routing even when the injected load is far away from the capacity of the system. In contrast, Web server farms that are allowed to reject requests can guarantee a high quality of service for every individual request stream even under relatively high injection rates.

In [104], members of Patras and Cyprus present the first complexity results concerning a particular selfish routing game. They prove that pure strategies in that game can be computed in polynomial time and also provide a polynomial time algorithm for mixed strategies. On the other hand, they show that it is #P-complete to compute the social cost of a particular mixed strategy, when the probabilities of the strategy are given. They also provide analytic evidence about strategies that achieve the worst-case social cost.

In [127], members of Patras and Cyprus extend the notion of Nash equilibria to that of approximate equilibria. They then present a novel probabilistic technique (balls-fusion) through which they derive a tight upper bound on the cost of anarchy of selfish routing games for such equilibria.

Load balancing

In [122] MPI studies the problem of assigning independent tasks to uniformly related machines when only limited online information is provided to the scheduler. This is an elementary scheduling problem having many applications in areas where the communication is costly and a scheduler should require the minimum online information in order to achieve a fixed performance, such as online load balancing and online allocation of decentralized resources.

The setting studied in [122] is a general framework whose special cases are the classical multiple-choice games for the assignment of independent tasks to identical machines. The latter case was the subject of intensive research for the last decade. The problem is intriguing in the sense that the natural extensions of the greedy oblivious schedulers, which are known to achieve near-optimal performance in the case of identical machines, are proved to perform quite poorly when distinct machine speeds come into play.

In particular, a rather surprising lower bound is presented stating that any oblivious scheduler that assigns an arbitrary number of tasks to n related machines would need $\Omega\left(\frac{\log n}{\log\log n}\right)$ polls of machine loads per task, in order to achieve a constant competitive ratio versus the optimum offline assignment of the same input sequence to these machines. On the other hand, it is claimed that the only missing information for an oblivious scheduler to perform asymptotically optimally, is the amount of tasks to be inserted into the system. To justify this, a new oblivious scheduler is provided that uses only $O(\log\log n)$ polls, along with the additional information of the size of the input sequence, in order to achieve a constant competitive ratio vs. the optimum offline assignment. The nature of the proposed scheduler is quite different from the existing schedulers for identical machines (which are mainly greedy schedulers) and is based on the concept of giving higher priority to the least favourable machines.

This oblivious scheduler is used as the core of an adaptive scheduler that does not demand the knowledge of the input sequence and yet achieves similar performance, within a constant factor.

In the same line of research, Paderborn mainly studied the distributed load-balancing problem. Distributed load-balancing is often done in two separated phases: *Flow Computation* and *Load Migration*. Concerning the first phase, a well known approach is based on diffusion algorithms. Several different diffusion schemes have previously been developed, especially for load-balancing

on homogeneous processor networks.

In [91], members of Paderborn generalize existing schemes in order to deal with heterogeneous networks. The generalized schemes may operate efficiently on networks where every processor can have arbitrary computing power, i.e. the load will be balanced proportionally to these weights. The balancing flow calculated by the schemes for homogeneous networks is minimal with regard to the l_2-norm and this is proved to be true for the generalized schemes, too.

[92] concentrates on heterogeneous networks where the capacities of the edges varying. It is shown that the convergence rate of diffusion algorithms can be improved using edge weighted graphs without deterioration of the quality of the flows. Moreover, in the same work, the authors consider common interconnection topologies and demonstrate how optimal edge weights can be calculated for the *First and Second Order Diffusion Schemes*. Using theoretical analysis and practical experiments it is shown that improvements can be achieved on selected networks.

[133] investigates how to schedule a flow on a synchronous distributed network. In this paper, members of Paderborn concentrate on the second phase of the two-step approach described above. The goal is to use the minimum number κ of rounds to reach the balanced state. This flow scheduling problem trivially leads to a formulation as a linear program with $\kappa(m+n)$ unknowns and $m+2\kappa n$ equations. (Such a system is solvable in $O(\kappa^5(m+n)^5)$ rounds using e.g. Karmarkar's algorithm.) The best distributed scheduling strategy so far needs at most $O(\sqrt{n})$ times the minimum number of rounds. It is shown that every distributed scheduling strategy requires at least $\frac{3}{2}$ times the minimum number of rounds. Furthermore, a distributed algorithm for scheduling flows on tree networks was given, which requires at most two times the optimal number of rounds.

In [143], members of Paderborn study the load-balancing problem with respect to a demand-driven parallelization of a ray tracing algorithm. Correctness and optimality of a perfect load-balancing algorithm for image space subdivision are proved and its exact message complexity is given. An integration of antialiasing into the load-balancing algorithm is proposed. A distributed object database allows rendering of complex scenes which cannot be stored in a single processor's memory. Each processor maintains a permanent subset of the object database as well as a cache for temporary storage of other objects. Using object bounding boxes and bounding hierarchy reduces the number of requests for missing data to a necessary minimum. The proposed parallelization is simple and robust. It should be easy to implement with any sequential ray tracer and any message-passing system.

In [157] evolving tree computations on circulant (rings with chords) and related graphs are analyzed. In an evolving α-ary tree computation, a complete tree grows level by level, i. e., every leaf generates α new nodes that become the new leaves. The load-balancing task is to spread the new nodes on a network of processors at the moment they are created in such a way that the accumulated number of nodes per processor, i. e., its load, is as close as possible to the average number of nodes per processor. As a single processor can hold many leaves, it has to handle many new tokens. Gao and Rosenberg introduced evolving computations and investigated the growth of complete binary trees on rings of processors. They showed that the so-called KS-regimen where one new node generated by a leaf remains at the processor where it has been created and the other one is transmitted to the right neighbouring processor behaves optimally in the course of long computations. In Paderborn's work, evolving computations are generalized to arbitrary trees and the regimen notion is extended. It is shown that even *any* regimen behaves optimally. For this purpose, they model the actual load distribution, the generation process, and the distribution regimen by formal infinite polynomials. Then it is shown that evaluating these polynomials at certain points leads to the analysis of these regimens on circulant and related graphs. It is shown that *any* regimen leads to a load distribution close to the optimal one.

An easy-to-use platform for efficient computing in distributed dynamic environments

In this context, Paderborn focuses 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, BSP algorithms for combinatorial problems are developed and analyzed and the impact of the model used on the design of these algorithms is investigated, on the other hand, a programming environment where one can actually execute BSP programs efficiently is implemented.

In the context of the first activity, an automatic configuration approach for implementing complex parallel BSP algorithms is investigated. [38] reports on the results of this approach. Here, a parallel algorithm is described by a sequence of instructions and of subproblems that have to be solved by other parallel algorithms called as subroutines, together with a mathematical description of its own running time. There may also be free algorithmic parameters, such as the degree of trees in the data structures used, that have an impact on the running time. As the running time of an algorithm depends on several machine parameters, on some fixed ones, and on the choice of the free algorithmic parameters and on the choice of the parallel subroutines for which the same statement applies in turn, the actual composition from all these ingredients of the parallel program for an actual parallel machine is a difficult task. Such a configuration system has been implemented using the Paderborn University BSP library. As an instructive example, members of Paderborn chose to implement a collection of sophisticated minimum spanning tree algorithms on two parallel machines, using the PUB library. One machine is a workstation cluster interconnected as a two-dimensional torus with SCI links. The other one is a workstation cluster where all processors are interconnected via a Fast Ethernet. The results of these implementations can be found in [38].

In the context of the second activity, Paderborn implemented a platform-independent programming environment continuing our previous work. Virtual processors were introduced and it was realized that their execution can be stopped on the current processor and resumed at a different one. This is implemented on a Linux workstation cluster. The actual release version 8.0 of the implementation of the Paderborn University BSP library can be inspected on the web page

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

A prototypical load-balancing mechanism has been integrated into this system.

WP 3: Production and Transportation Planning

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.

Below we describe the work done in WP3 in the second year of ALCOM-FT. We distinguish between research that is targeted at specific problems in this area and research that is more generic in the sense that it aims at improving algorithms that can be used to tackle different problems in this field. We start with the problem specific research.

Planning, Rostering, and Scheduling problems

In [41], Aarhus investigated the seat reservation problem. This is an on-line problem in which passengers are assigned to a seat in the train for the part of the trip that they want to travel. The objective is to maximize the sum of the ticket prizes, but only fair algorithms are allowed, which cannot refuse a passenger if a seat is available. Several on-line algorithms have been studied with respect to their competitive ratio. If a passenger must keep his seat for the whole trip, then the competitive ratio is $1/2$. The competitive ratio can be improved by allowing seat changes: if s seat changes per passenger are allowed, then an algorithm is possible with competitive ratio $(s+1)/(s+2)$. Furthermore, if one seat change is allowed, then an algorithm is presented that needs at most twice the number of seats that is required by an optimal off-line algorithm to seat everybody. In [94, 95], the competitive ratio is studied that is achievable for a scheduling problem with two uniform machines with the objective of minimizing the makespan. The problem is not fully on-line, but semi on-line, since the jobs are supposed to arrive in order of non-increasing

processing times. In [94] the competitive ratio is analyzed as a function of the ratio q of the machine speeds. We analyze the LPT algorithm, which always assigns the longest job to the machine that becomes available first, and several adaptations to this algorithm. In [95] we make the additional assumption that idle time is allowed. We present optimal algorithms for all values of q, where quite surprisingly for $q > 2$ it may be necessary to leave a machine idle if there is work available.

Paderborn continued its work on planning problems. One of these concerns a model that reflects real-world requirements within a routing and scheduling application: Together with industrial partners, we integrate yield management and fleet assignment facilities. The aim of the integration is to achieve a higher revenue based on realistic prognostic models used by airline companies. Encouraging results have been presented already at this year's meeting of the *AGIFORS*[1] *Study Group on Reservation and Yield-Management*. A second project in the airline industry concerns the problem of integrating the fleet assignment and aircraft rotation phase, e.g. the assignment of fleet to the flight network with respect to individual restrictions like maintenance intervals, or time buffers for delay minimization. Furthermore, Paderborn continued its research on complex nurse rostering problems. In close co-operation with industrial partners, refinements on the models developed during the last year have been performed and introduced into the software. Since end-users do not need proven optimal solutions, but require good solutions quickly, we currently focus on some metaheuristic approaches for the nurse rostering problem.

In [153], a branch-and-bound approach is presented for the Capacitated Network Design Problem, which boils down to finding an optimal subset of the edges of a network such that a given set of commodities can be transported at minimum total cost. The focus is on tightening strategies such as variable fixing and local cuts that can be applied in every search node. Different variable fixing algorithms based on Lagrangian relaxations are evaluated solitarily and in combined versions. Moreover, cardinality cuts are developed for the problem and their usefulness is evaluated empirically by numerous tests.

Utrecht has continued its work on scheduling problems. In [2], a problem of accepting or rejecting jobs is studied, where jobs have stochastic instead of deterministic processing times. If a job is accepted, then it must be completed at a given due date with a certain given threshold probability. Four problem classes with stochastic processing times are considered. In the first three classes the jobs have processing times that follow: (i) a gamma distribution with parameters α_j and β, where β is common to all jobs; (ii) a negative binomial distribution with parameters s_j and p, where p is the same for each job; (iii) a normal distribution with parameters μ_j and σ_j^2. The fourth class has equally perturbed processing times, that is, the processing time consist of a deterministic part and a random component that is independently, identically distributed for each job. While the first two classes of problems can be solved by the same algorithm as the deterministic case, a more intricate dynamic programming algorithm is needed for the other two classes, which requires pseudo-polynomial time in case of normally distributed processing times. In [3], a variant of the two machine open shop problem is considered, where the question whether each machine can be ready before a given deadline is resolved. In [71], a tabu-search algorithm for a purchase problem is developed, where the goal is to minimize the purchase and holding cost; one of the complicating (and more realistic) features was the presence of quantity discounts. This is one of the components of supply chain optimization, in which the supplier and the manufacturer jointly optimize their planning. In [72, 73], the minimum test set problem is studied. In this problem one encounters a set of items that have to be uniquely identified; for the identification process a number of tests is present. Each test contains a number of items; it distinguishes between items i and j if exactly one of these is contained in the test. The objective is to select a minimum test set. In [73] a branch-and-bound framework consisting of several branching rules and lower bounds is presented. In [72] it is shown that the greedy algorithm is best from a worst-case point of view. We further present and analyze a number of heuristics for the special case in which each test contains at most two items, and we show that the problem is APX-hard.

[1] AGIFORS: Airline Group of the International Federation of Operational Research

Methodological research

In this part we describe the research that has been conducted on algorithms that are used to solve planning problems.

Cologne continued its development of the package SCIL (Symbolic Constraints in Integer Linear Programming), which will become very useful to solve the problems encountered in WP3. A prototype has been released recently. We refer to the section on WP4 for a full description.

Paderborn has derived several nice results in the area of constraint satisfaction. When applying constraint satisfaction one has to explore the search space and discard infeasible solutions. If there are many symmetric, infeasible solutions, then the search algorithm will have to invest time on each one of them before it can discard them, even though they are in fact equal. In [98], a method is presented to check during the search procedure whether a new point is symmetric or dominated by a solution that has been discarded before. Numerical experiments indicate that this method speeds up the search process considerably. Another method to speed up constraint satisfaction is to reduce the number of nodes that have to be explored in search tree; this is presented in [152]. The key idea here is to use redundant constraints, which are propagated heuristically. The effectiveness of the approach is demonstrated on the so-called Social Golfer Problem. In [97], a branch-and-bound algorithm is presented for the maximum clique problem, which is one of the basic problems that is often encountered as a subproblem. Cost based filtering techniques are introduced for the so-called candidate set (i.e. a set of nodes that can possibly extend the clique in the current choice point). Additionally, a taxonomy of upper bounds for maximum clique is presented. Analytical results prove that the cost based filtering is in a sense as tight as most of these well-known bounds for the maximum clique problem. Experiments show that the combination of cost based filtering and vertex coloring bounds outperforms other approaches as well as approaches that only apply either of these techniques.

Patras in [151] presents a space reduction approach by which the time needed to compute a shortest path in a graph is reduced considerably. This research was motivated by the problem of retrieving information concerning the time table of the German railways, but there are many other possible applications.

Saarbrücken [33] studied the problem of how to position convex polytopes such that they do not overlap. Utrecht [123] presented several upper and lower bounds to compute the tree-width of a graph. The knowledge of the tree-width is necessary to find a good tree decomposition of the input graph, which may lead to a faster algorithm for the corresponding graph problem.

WP 4: Generic Methods

During this second year of the Project, a total of 106 reports were produced within WP4. Of these, nearly half were assigned jointly with other work packages, notably with WP2 (26), demonstrating the intimate interweaving among work packages. More than half of the reports have appeared or are about to appear in journals or in conference and workshop proceedings. We will summarise below some of the identifiable themes represented by the year's reporting.

Online Algorithms

An important issue when working with on-line algorithms is how to evaluate the algorithms. The standard measure for the quality of on-line algorithms is the competitive ratio, which roughly speaking is the worst case ratio of the performance of the on-line algorithm to the performance of an optimal off-line algorithm. However, sometimes this measure gives results that are overly negative compared to empirical results, and fails to distinguish algorithms that are known to perform very differently in practice.

This is especially true for the paging problem, where the competitive ratio of almost any reasonable algorithm equals the size of the cache. The problem is that using the competitive ratio no structure of the sequences is assumed and, for sequences that do not exhibit locality of reference, a cache is of no use. In [5] we devise a simple way of modeling locality of reference. In

this model we are able to evaluate the algorithms using a very natural measure, the fault rate. This gives results closer to reality than those obtained for the competitive ratio.

In [93] we study the bin packing problem, where a fixed number of bins are given, and the aim is to pack as many items as possible in the bins. It has been shown earlier that in the general case, no algorithm for this problem has a constant competitive ratio, even in the case of identical bins. However, if we study only item sequences that can be packed completely by an optimal off-line algorithm, a whole class of very natural algorithms have competitive ratios of at least $\frac{1}{2}$, even in the case where the bins may have different sizes.

The idea of considering only sequences for which the amount of resources available (like the bins in the bin packing problem) suffices for an optimal off-line algorithm has earlier been generalized to considering sequences for which α times the amount of resources available ($\alpha \geq 1$) suffices for an optimal off-line algorithm. In [40], the case $\alpha < 1$ is studied, i.e., the amount of resources available is more than sufficient for an optimal off-line algorithm. The competitive ratio as a function of α ($\alpha < 1$) is studied for various problems, e.g., bin packing and scheduling.

Approximation methods

Our goal is to find fast approximation algorithms with constant worst case ratio. One of the ways of achieving this discussed in [100] is by allowing our approximation algorithms access to a certain advice called a d-octopus of the graph. Among others, we consider the k-clustering problem which appears quite naturally in the context of transport, warehousing and networking, the bandwidth minimization problem which plays an important role in computational linear algebra, and radio labeling problems. Another approach to obtaining polynomial-time algorithms to approximate a minimum cocolouring on graphs, partially ordered sets and sequences is considered in [101]. In particular, we obtain an efficient algorithm to approximate within a factor of 1.71 a minimum partition of a partially ordered set into chains and antichains, and a minimum partition of a sequence into increasing and decreasing subsequences.

The time-dependent orienteering problem is dual to the time-dependent traveling salesman problem. It consists of visiting a maximum number of sites within a given deadline. The traveling time between two sites is in general dependent on the starting time. Some greedy algorithms for the time-dependent orienteering problem with good worst case ratio were obtained in [102]. For any $\epsilon > 0$, we provide a $(2 + \epsilon)$-approximation algorithm for the time-dependent orienteering problem which runs in polynomial time if the ratio between the maximum and minimum traveling time between any two sites is constant. No prior upper approximation bounds were known for this time-dependent problem.

When the linkage structure of communication networks is modeled by a graph, the diameter of the graph corresponds to the maximum number of links over which a message between two nodes must travel. In cases where the number of links in a path is roughly proportional to the time delay or signal degradation encountered by messages sent along the path, the diameter is involved in the complexity analysis for the performance of the networks. The oriented diameter of an (undirected) graph G is the smallest diameter among all the diameters of strongly connected orientations of G. In [103] we study algorithmic aspects of determining the oriented diameter of a chordal graph. We give a linear time algorithm which, for a given chordal graph G, either concludes that there is no strongly connected orientation of G, or finds a strongly connected orientation of G with diameter at most twice the diameter of G plus one. Moreover, we prove that the corresponding decision problem remains NP-complete even when restricted to a small subclass of chordal graphs called split graphs. Finally we show that unless P=NP, there is neither a polynomial-time absolute approximation algorithm nor an α-approximation (for every $\alpha < \frac{3}{2}$) algorithm computing the oriented diameter of a chordal graph.

Graph Algorithms

A bisection of a graph is the partition of its vertices into two equal subsets, and the size of a bisection is the number of edges between the two subsets. A cubic graph has every vertex of

degree three. Two randomized algorithms to compute the bisection width of random cubic graphs with n vertices are presented in [80], giving asymptotic upper bounds for the bisection width of $0.174039n$ and $0.174501n$ respectively. We also obtain an asymptotic lower bound for the size of the max bisection of a random cubic graph with n vertices of $1.325961n$ and $1.325499n$. The tight analysis is based on the differential equation method. The results give an asymptotically almost surely width size for the bisection on random cubic graphs. The technique differs from the typical differential equation method in that we need to use a deprioritized algorithm. It also yields a random constant-approximation to MAX CUT on a random cubic graph.

In [51] we consider the problem of colouring a planar graph with the minimum number of colours such that each colour class avoids one or more forbidden graphs as subgraphs. We perform a detailed study of the computational complexity of this problem. We present a complete picture for the case with a single forbidden connected (induced or non-induced) subgraph. The 2-colouring problem is NP-hard if the forbidden subgraph is a tree with at least two edges, and it is polynomially solvable in all other cases. The 3-colouring problem is NP-hard if the forbidden subgraph is a path, and it is polynomially solvable in all other cases. We also derive results for several forbidden sets of cycles. In [52] we continue the study of colouring problems with forbidden subgraphs. A subcolouring is a vertex colouring of a graph in which every colour class induces a disjoint union of cliques. We derive a number of results on the combinatorics, the algorithmics, and the complexity of subcolourings. On the negative side, we prove that 2-subcolouring is NP-hard for comparability graphs, and that 3-subcolouring is NP-hard for AT-free graphs and for complements of planar graphs. On the positive side, we derive polynomial time algorithms for 2-subcolouring of complements of planar graphs, and for r-subcolouring of interval and of permutation graphs. Moreover, we prove asymptotically best possible upper bounds on the subchromatic number of interval graphs, chordal graphs, and permutation graphs in terms of the number of vertices.

Research continued on the efficiency of computing graph and network problems related to the notion of treewidth. Computational experiments on computing the treewidth of networks, coming from applications on decision support systems and telecommunication, are reported in [123], with comparisons made between existing and new heuristics for this problem. In [37] it is shown how in several cases, the modular decomposition of a network can be used to solve the treewidth problem and the related problem of minimum fill-in. A new notion of the cost of a tree decomposition is introduced [36], that captures better than the notion of treewidth the time needed when using tree decomposition to solve specific problems, and algorithms to compute tree decompositions of minimum cost are given.

Property testing

A *property testing algorithm* has just to distinguish between the case when a given object has a predetermined property and the case when the object is far away from any object having the property. For example, for a geometric object such as a set of points or a geometric graph, we might consider properties such as whether the set is convex or whether the graph is a Euclidean Minimum Spanning Tree, or whether they are far from any object that has this property.

When we want to describe a property testing problem we have to specify a *distance measure* and the type of query used by the algorithm. The distance between objects is typically the fraction of 'atomic items' on which the two objects (or their representations) differ. For example, in the case of point sets this is the fraction of points that are in one point set but not in the other. An object is ϵ-far from a property if it has a distance of more than ϵ from any object having the property. In the 'standard testing model' we allow only the most basic queries about the object. For example, when the considered object is a point set then we may ask queries of the form 'What is the position of the i-th point of the set ?'.

The goal of a property testing algorithm is to decide whether the object has the property or is far from the property much more quickly than would be possible for the corresponding (exact) decision problem. Often it is possible to test a property in time independent of the object description size (but dependent on the distance parameter ϵ).

However, there are many geometric properties that cannot be tested efficiently in the standard

model. This might change when the algorithm is allowed to ask more powerful queries than in the standard model, for example queries of the form: 'What is the i-th point in the query range $\mathcal{R}$?', where $\mathcal{R}$ is typically an axis-parallel hyperrectangle or a simplex. Such queries are efficiently supported by basic geometric data structures such as range trees or partition trees. In [68] we introduce range queries for property testing and show that some properties of point sets (convex position, and labeling and clustering properties) can be tested much more efficiently if the testing algorithm is allowed to use such range queries.

Dynamic Graph and Hypergraph Algorithms

Part of our effort was devoted to studying dynamic path problems on directed graphs. In particular, we have considered the problem of maintaining all-pairs shortest paths in a weighted directed graph.

In [78], we have shown how to obtain effective query/update trade-offs for dynamic all-pairs shortest paths, by introducing two new families of algorithms. Given a weighted directed graph with n vertices where each edge can attain S different real values, algorithms in the first family achieve an update bound of $\widetilde{O}(Skn^2)$ and a query bound of $\widetilde{O}(n/k)$, and improve over the best known update bounds for any k in the range $(n/S)^{1/3} \leq k < (n/S)^{1/2}$. Algorithms in the second family achieve an update bound of $\widetilde{O}\left(Skn^2\right)$ and a query bound of $\widetilde{O}(n^2/k^2)$, and are competitive with the best known update bounds (first family included) for any k in the range $(n/S)^{1/6} \leq k < (n/S)^{1/3}$.

In a more recent paper [77], we have studied novel combinatorial properties of graphs and we have devised a completely new approach to dynamic all-pairs shortest paths based on these properties. Our approach yields a fully dynamic algorithm for general directed graphs with real edge weights that supports any sequence of operations in $O(n^2 \log n)$ amortized time per update and unit worst-case time per distance query, where n is the number of vertices. We remark that this is the first algorithm that solves the dynamic all-pairs shortest paths problem in its full generality. Furthermore, our bounds improve substantially over previous results.

Finally in [18], we have addressed the problem of designing efficient semidynamic algorithms for hypergraph shortest path problems. In particular, in the decremental case, the proposed approach provides the best known performance for such problem.

Data Structures

Inspired by recent work of Buhrman et al. on space-efficient randomized analogues of bit vectors, we consider dictionaries that perform lookups by probing a *single word* of memory, knowing only the size of the data structure. We describe a randomized dictionary where a lookup returns the correct answer with probability $1 - \epsilon$, and otherwise returns "don't know". We show that there exists a (non-explicit) one-probe dictionary with success probability $1 - \epsilon$ for any $\epsilon > 0$. It uses $O(n \log \frac{2u}{n})$ words of memory, where u is the size of the universe and n is the number of elements in the dictionary. An explicit construction using $n2^{O((\log \log u)^3)}$ words of memory is also given. In contrast, for deterministic dictionaries, space proportional to u is needed.

Characterizing and finding regularities in strings are important problems in e.g. text compression and computational biology. Apostolico and Preparata introduced the *string statistics problem*, which consists of preprocessing a string such that given a query pattern, the maximum number of non-overlapping occurrences of the query pattern can be reported efficiently. In [45] we describe an algorithm for this problem which uses time $O(n \log n)$, where n is the length of the string, to build a data structure such that queries can be answered in time $O(m)$, where m is the length of the pattern. This improves on the algorithm of Apostolico and Preparata by a logarithmic factor.

We have studied the problem of constructing (static) compact data structures for representing information about all-pairs shortest paths in a weighted directed graph. In particular, for a directed graph G we have considered queries of the form: "What is the shortest path distance from vertex x to vertex y in G avoiding a failed link (u, v), and what edge leaving x should we use to get on a such a shortest path?". In [79], we have shown that an oracle for such queries can be

stored in $O(n^2 \log n)$ space with a query time of $O(\log n)$. No non-trivial solution was known for this problem.

Combinatorial Methods

For most algorithms, there is a discrepancy between the worst-case and the average-case behaviour, and both quantities convey very useful information. A new kind of complexity analysis, called *smoothed analysis*, was introduced last year by Spielman et al. This is intermediate between average-case analysis and worst-case analysis. It consists in studying the performance of a given algorithm on a perturbation of the worst-case instance(s). For worst-case analysis, one is interested in the quantity $\max_x C(x)$, the highest cost of the algorithm, for inputs x in a given set of data. For average-case analysis, one is interested in the quantity $\mathrm{E}[C(x)]$, the average cost of the algorithm over all the inputs x in a given set of data. For the so-called smoothed analysis, the quantity of interest is

$$\max_x \mathrm{E}_\epsilon[C(x+\epsilon)]\,,$$

when $x+\epsilon$ is an "ϵ-perturbation of x".

The main aim of smoothed analysis is to know if worst instances are rather isolated in the "complexity landscape" of the algorithm. For example, Spielman proved recently that the simplex algorithm in linear programming has polynomial "smoothed complexity", despite an exponential worst-case behaviour.

The concept of "discrete smoothed analysis" for algorithms running on discrete data structures is introduced in [27]. As there is no canonical definition for the notion of perturbation of the instances in this case, we consider here two natural (practically speaking) models of limited randomness: the first one corresponds to a *partial permutation* (of the instances) and the second corresponds to a *partial randomization* (of some bits of the instances). Classical algorithms such as quicksort and shortest paths in weighted graphs are successfully analysed under these models.

The report [39] addresses fundamental problems regarding the statistical occurrence of patterns in texts obeying quite general randomness models. The shape of the patterns and the characteristics of the information source are both more powerful than what has been considered so far, e.g. sources may have unbounded (non Markovian) correlations. This work complements research done in Year 1 and is motivated by the statistical analysis of large sequences as may appear in bio-informatics and intrusion detection for example.

The two reports [25, 28] deal with random walks with positivity constraints. The problems originated from queueing theory, and have applications in combinatorics and algorithmics. Indeed, it was recently observed that the shape of a generating tree (an important structure for uniform random generation of combinatorial objects) can be predicted if a given set of random walks can be enumerated. Closed-form formulas and asymptotics for various different parameters are given using the so-called "kernel-method".

Report [150] also has a "random combinatorics" flavour. Balls-in-urns problems are here revisited: a constraint of adjacency is taken into account. It corresponds to overlapping allocation problems of files for disks and drums. Probabilistic/analytic methods (involving the Lambert function) allow the amount of overlapping to be bounded.

A wide synthesis of combinatorial enumeration is presented in [99]. It develops in nearly 200 pages the basics of combinatorial enumeration through a "symbolic" approach that revolves around generating functions. The major objects of interest here are words, trees, graphs, and permutations, which surface recurrently in all areas of discrete mathematics. It is largely oriented towards applications of combinatorial enumeration to random discrete structures and discrete mathematics models, as they appear in various branches of science, like statistical physics, computational biology, probability theory, and, last but not least, computer science and the analysis of algorithms.

SCIL (Symbolic Constraints for Integer Linear Programming)

The software framework SCIL is a joint project by the MPI and Cologne sites together with Alexander Bockmayr (Nancy) and Thomas Kasper (SAP AG). The aim is to design a high-level description language for integer linear programming based algorithms for combinatorial optimization. The system is inspired by the constraint programming concept of symbolic constraints. In particular we want to create an easy-to-use system that introduces symbolic constraints in branch-and-cut-and-price (BCP) algorithms and supports flexible and compact modeling.

The development of SCIL is still an ongoing process. During the second year the core modeling framework reached a stable state (deliverable D20). It is available as a `C++` class library built on top of LEDA and ABACUS. The ALCOM-FT report [7] describes the concepts and usage of the system. For the current state of the implementation and more information, refer to the web-site

`http://www.mpi-sb.mpg.de/SCIL/index.html`

Two new directions to enhance SCIL have been identified. First, we want to provide a BCP-system, that can guarantee termination and optimality although the underlying LP-solvers use inexact floating point arithmetic. We hope to attract less experienced users who are not aware of numerical problems within LP-solvers and BCP-systems. Second, we want to extend SCIL by decomposition methods like Lagrangean relaxation and approximate LP-solving, for users who want to find good solutions fast without asking for optimality. A first version of the relaxation code has been implemented but is still in prototype.

WP 5: Experimental Algorithmics

The 20 ALCOM-FT reports [4,5,20,48,53,58,62,76,84,85,88,121,139,144–148,151,158] have been classified as WP5. Until now, 11 of these have appeared in conference proceedings or journals, or are scheduled to do so soon. In most of our groups it has become more or less standard to augment a significant fraction of our theoretical work with experiments. Indeed, it is quite likely that experimental aspects appear in more of the ALCOM reports than are listed here. It also makes little sense to enumerate all the specific topics that have an algorithm engineering aspect. Rather we highlight a few specific areas where several papers have appeared, omitting others that have only a few papers this year or that continue work explained in the report from last year.

Related to this work package, Rasmus Pagh and Fleming Rodler received the best student paper award at the ESA 2001 conference for a practical hash table data structure that has worst case constant access time, and Kurt Mehlhorn gave a plenary invited talk on LEDA at SODA 2002, the largest algorithms conference world wide.

Below we report in more detail on an ALCOM-FT School on Algorithm Engineering, on a package for computing optimal Steiner trees, dynamic graph algorithms, cache oblivious algorithms, algorithm visualization, large scale rendering systems, and on uniform random generation of combinatorial objects.

ALCOM-FT School on Algorithm Engineering

On September 10-12, 2001, Rome organized a "School on Algorithm Engineering", sponsored by ALCOM-FT and co-organized by Giuseppe F. Italiano (University of Rome "Tor Vergata") and by Giovanni Rinaldi (IASI-CNR). The School was devoted to PhD students and researchers interested in algorithm engineering and experimentation of algorithms, and took place in Rome at IASI-CNR, Viale Manzoni 30. The format of the School was the following: there were lectures delivered by the invited speakers, and there were also time for talks and demos, organized in an informal and cozy atmosphere, in order to leave plenty of opportunities for discussions and interactions among participants.

The invited speakers were:

1. Michael Jünger, Universität zu Köln: "Branch and Cut Algorithms for Combinatorial Optimization"

2. Kurt Mehlhorn, Max Planck Institut für Informatik, Saarbrücken "Software Libraries for Efficient Algorithms"

3. Bernard Moret, University of New Mexico: "Engineering Algorithms for Speed: Principles, Challenges, and Case Studies"

The school was quite successful: the participants were 59, coming from 12 different countries (Austria, Belgium, Denmark, France, Germany, Greece, Hungary, Italy, Spain, Switzerland, UK, USA).

Steiner trees

Reports [145–147] are part of an ongoing effort to write a package that can solve large real-world Steiner tree problems optimally. Steiner trees are one of the most widely studied network design problems with applications, for example, in VLSI design. Our system is among the world leading systems in the sense that it is the fastest program for most instances from the well-established benchmark library SteinLib, and that for some large instances it is the only program finding optimal solutions. This system integrates many preprocessing and bounding techniques, two of which we describe here.

- Preprocessing methods are a key ingredient of the most successful algorithms for the Steiner problem. Whereas classical methods just considered single vertices or edges as objects for reduction tests, recent and more sophisticated tests extend the scope of inspection to more general patterns. In [145], we present such an extended reduction test, which generalizes different approaches in the literature. We introduce the new approach of combining alternative- and bound-based methods, which substantially improves the impact of the tests. The experimental results presented there show a large improvement over previous methods using the idea of extension.

- Partitioning is one of the basic ideas for designing efficient algorithms, but on NP-hard problems like the Steiner problem, straightforward application of the classical paradigms for exploiting this idea rarely leads to empirically successful algorithms. In [147], we present a new approach that is based on vertex separators. We show several contexts in which this approach can be used profitably. Our approach is new in the sense that it uses partitioning to design reduction methods. We introduce two such methods and show their impact empirically.

Dynamic graph algorithms

In [58], we conduct an extensive empirical study on several algorithms for maintaining minimum spanning trees in dynamic graphs. In particular, we have implemented and tested a variant of the polylogarithmic algorithm by Holm et al., sparsification on top of Frederickson's algorithm, and compared them to other (less sophisticated) dynamic algorithms. In our experiments, we considered as test sets several random, semi-random and worst-case inputs previously considered in the literature.

A detailed survey of experimental studies on dynamic graph algorithms was prepared in [158].

Cache oblivious algorithms

Recently, the concept of *cache oblivious* algorithms has been introduced by Frigo et al. as an elegant way of dealing with hierarchical memory systems (more details on the model can be found in the section on Work Package 1). It is a theoretically very appealing model, with probable practical significance, but the latter point largely remains to be verified.

In [48], we propose a simple cache oblivious analog of standard (cache aware) B-trees, which can be implemented implicitly as a single array of data elements without the use of pointers. We also try to assess systematically the impact of the memory layout of search trees by experimentally comparing the efficiency of the cache-oblivious van Emde Boas layout with a cache-aware layout based on multiway trees, and with classical layouts such as breath first search, depth first search, and in-order. Our empirical results indicate that the nice theoretical properties of cache oblivious search trees actually do carry over into practice. Another observation is that the effects from the saving in space and increase in fanout caused by implicitness are notable.

Further papers on memory hierarchies are [5, 121, 148].

Algorithm visualization

Part of our effort was devoted to studying techniques for the visualization of algorithms [76]. In particular, we have considered two of the most popular visualization specification techniques, i.e., the "interesting events" and the "state mapping" methods, providing the first direct comparison of these two approaches and identifying scenarios where one might be preferable to the other. We have based our study on concrete examples realized in the algorithm visualization system Leonardo, developed at our site, and in the algorithm visualization system Polka, developed at the GVU Center - Georgia Institute of Technology.

Rendering of highly complex scenes

Real-time rendering of large 3D scenes using advanced techniques like radiosity and ray-tracing is perhaps the most dynamically growing area in terms of the consumption of computing resources. The Paderborn side works on new data structures and efficient distributed implementations in this area.

In [144] we present the development and testing of a prototype of an e-commerce system which simulates realistic lighting of large scenes on high performance parallel computers. The system, although tailored to the needs of a specific application, is very generic and exhibits metacomputing features: 1. The access to high performance computers is fully transparent to the user; 2. The modular architecture of the system allows one to dynamically add or remove computing resources in geographically different computing centers. The prototype of the proposed system was evaluated in the industrial contexts of architectural visualization and film production. The paper summarizes scientific and technical problems which arose during the project as well as their solutions and engineering decisions.

In [121] we present improved algorithms of our preceding work about the rendering of highly complex 3D scenes of arbitrary topology. The specific feature of our method is that it makes possible an interactive navigation of scenes consisting of more than 16 GB of polygon data (95 million polygons) that cannot be stored in the main memory, but only on a local or remote hard disk. For the computation of an approximate image of the scene, a sampling technique is used. In the preprocessing, a so called sample tree is built whose nodes contain randomly selected polygons from the scene. This tree uses only space linear in the number of polygons. In order to produce an image of the scene, the tree is traversed and polygons stored in the visited nodes are rendered. During the interactive walkthrough, parts of the sample tree are loaded from local or remote hard disk.

Visualizing is a method used to help the experiencing and understanding of causal cohesions in simulation processes. For this purpose, tools for visualizing are already implemented in prevalent simulation systems. The user creates his simulation model and generates a 3-dimensional (2.5-dimensional) visualization by means of the simulation system. This helps the examination of the process which makes it easier for the viewer to "understand" it. Simulation tools usually only provide the opportunity for a unidirectional visualizing. In a 3-dimensional surrounding the viewer cannot implement an interaction with the simulation while the system is running. In [139] we present a distributed system consisting of a commercial manufacturing simulation tool, a coupling module, and a walkthrough system. The distributed system in conjunction with

the coupling module guarantees generality and a wide field of applications of the walkthrough system. Further it guarantees flexibility and selection of the specialized graphics hardware for the walkthrough system.

Uniform random generation of combinatorial objects

Uniform random generation allows one to study experimentally the behavior and complexity of algorithms by running them with thousands of randomly generated inputs.

The paper [85] deals with uniform random generation of combinatorial objects (of size $n(1+\epsilon)$, for a given ϵ). Most of the known algorithms for uniform random generation (of objects of size n exactly) have an average complexity of $O(n^2)$. This paper, by allowing a "relaxation" factor ϵ for the size, gives a general way to get very efficient algorithms ($O(n)$ in time and $O(1)$ as extra storage in memory).

This is in fact a rejection algorithms for which the number of rejections is minimized, thanks to a mathematical analysis of the underlying combinatorial structure. Another advantage of this process is that an automatization (via computer algebra systems) is possible.

WP 6: Project Management, Dissemination, Evaluation

Project Management

The Consortium Board, consisting of the scientific leaders of each of the eleven participating sites and chaired by the Coordinator, has carried the overall responsibility for the management of the project. In frequent contact with the members of the Board, 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.

The work at each site has been organized by the site leader in question, with the help of one or more assistants. The coordination of the activities within each of the six work packages has been the responsibility of the Work Package Leader who has also performed the reporting on the work done in the work package.

The second ALCOM-FT meeting of the Consortium Board was held on September 14, 2001 in Rome. The Board found the project to be in good shape. In total, slightly less than the planned one third of the funds had been spent, and the work had proceeded as planned. However, concerns were raised about the problem of hiring personnel for the more programming-intensive deliverables.

This concern later turned out to be well-founded. In year two, the work on several deliverables has been delayed due to problems with filling positions for programmers. This applies to deliverables D14, D15, D19, and D21, for which the related programmer positions took significantly more time than anticipated to fill. These deliverables are therefore somewhat delayed, but are expected to be completed during the fall of 2002 (D19 possibly even later).

One further outcome of the Board meeting was the decision to include the group of Marios Mavronicolas at the University of Cyprus in the ALCOM-FT consortium. The strengths of this group in the areas of distributed computing and networks fit nicely into the project, in particular with the Work Packages 2 and 4. The group joined the consortium on March 1, 2002.

The first annual ALCOM-FT review meeting and workshop was held in Rome, September 13–15, 2001. The review meeting featured a presentation of the work done in each of the work packages of the project, and at the ensuing scientific workshop, each site presented details of selected work from the first year of the project.

As a further means in the steering of the project, an external review of a few key deliverables has been initiated.

Dissemination

The principal means of dissemination for the ALCOM-FT project is scientific reports describing the research done. A total of 158 reports have been produced during the second year of the project. They have all appeared in the ALCOM-FT Technical Report Series, which is available online.

During the second year of the project, two Summer Schools were organized. The first was held at MPI in Saarbrücken, September 2001, and had the title *2nd Max-Planck Advanced Course on the Foundations of Computer Science*. The second was held in Rome, September 2002, and had the title *ALCOM-FT Summer School on Algorithm Engineering*.

The ALCOM-FT website (`http://www.brics.dk/ALCOM-FT`) has been continuously updated to contain the latest deliverables produced, including the 158 new reports in the ALCOM-FT Technical Report Series. Several of the deliverables have separate websites, which can be reached through the ALCOM-FT website.

During the second year of the project, the conferences STOC, ICALP, and SPAA were held in Crete, Greece, organized by ALCOM-FT site CTI. The conferences ESA, WAE, and WABI were held in Aarhus, Denmark, organized by ALCOM-FT site Aarhus. These are among the most important conferences within algorithmics world-wide. Additionally, many smaller conferences and workshops with subjects relevant for the project were arranged by various ALCOM-FT sites.

As a means to increase the awareness of algorithmics in research communities outside ALCOM-FT, we have set up two algorithmic competitions based on finding worst case instances for graph algorithms from the LEDA library. The prize is a stay of one week at an ALCOM-FT site of choice, for a total cost of up to EUR 1000, including travel. Further details of the competitions can be found at `http://www.mpi-sb.mpg.de/~schaefer/MLLB/index.html`. The competitions are open to everyone interested.

Dissemination of algorithmic awareness has also been fostered by talks given by partners to audiences in industry. As examples, Michael Jünger of the Cologne site talked on *Automatic Layout of Diagrams* at the Chamber of Industry and Commerce of Cologne, José Balcazar of the Barcelona site talked on *Support vector machines: survey and new algorithms* at Mannes Technology and at Hewlett-Packard labs, and the Paderborn site has established a forum called *Industrie trifft Informatik* (Industry meets Computer Science), where university and business cultivate their dialog in the field of information technologies and where possibilities for new cooperations are launched (see `http://www.uni-paderborn.de/cs/kooperation/iti.html` for further details).

On the other hand, we have decided to cancel deliverable D14, a proposed website with PR-pages for algorithmics. Reviewing the currently existing resources on the web, we have found that our efforts in that direction would be a duplication of efforts already done elsewhere (see e.g. `http://www.cs.sunysb.edu/~algorith/` for an example of existing material).

Evaluation

Overall, the ALCOM-FT project is in good shape. The work is basically progressing according to plans and the productivity continues to be very high.

The international standing of ALCOM-FT is illustrated by the fact that a majority of the total of 158 reports have already been published (or been accepted for publication) in the scientific community via conferences and journals, including the following key conferences: ALENEX, ESA, ICALP, MFCS, SPAA, SODA, STACS, STOC, SWAT, and WAE.

Some problems with vacant programming positions have been encountered, and these have caused four deliverables to be slightly delayed. The situation is under control, as most positions have been filled.

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 primarily 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 11 associate/assistant professors and 9 Ph.D. students. Rolf Fagerberg and Anna Östlin are paid part-time by ALCOM-FT funds. Rolf Fagerberg is responsible for running the day-to-day business of the project.

During the first year of ALCOM-FT, Riko Jacob completed his Ph.D. with the thesis *Dynamic Planar Convex Hull* and Lars Jacobsen completed his Ph.D. with the thesis *Search Trees with Local Rules*. Visitors to the site include L. Epstein (The Interdisciplinary Center, Herzliya), S. Srinivasa Rao (Leicester) and V. Arvind (IMSc, Chennai)

Peter Bro Miltersen served in the program committee of Computational Complexity 01, ESA 01, RANDOM 01, FOCS 01, and FSTTCS 01. Gerth S. Brodal served in the program committee of ESA 01. Erik Meineche Schmidt was on the program committee of ICALP 01. Sven Skyum was a member of an international committee formed by the National Agency for Higher Education (Högskolaverket) in Sweden to evaluate the standard of higher education in Computer Science in Sweden.

Members of the site have participated in most of the important algorithm conferences, including ESA, FOCS, ICALP, SODA, STOC, and SWAT. In August 2001, the site hosted the event *ALGO 2001*, which comprised the three conferences *9th Annual European Symposium on Algorithms (ESA)*, *5th Workshop on Algorithm Engineering (WAE)*, and *1st Workshop on Algorithms in BioInformatics (WABI)*.

The 23 ALCOM-FT reports [5, 6, 40–42, 44–50, 65, 70, 93–95, 105, 112, 128, 134, 140, 155] 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 in the UPC. This second period, the site is involved in Work Packages 1, 2 and 4. In total, the groups currently comprise 10 associate/assistant/full professors and 4 Ph.D. students. Several students have been paid part-time by ALCOM-FT funds, to do tasks of implementation.

During the second year of ALCOM-FT, Albert Atserias completed his Ph.D. with the thesis *The complexity of resource-bounded propositional proofs*.

Visitors to the site include , Philippe Flajolet, Paul Spirakis, Daniel Panario, Alfredo Viola, Jarik Nesetril, Eythan Levy, Elena Prieto and Nick Wormald. Dimitrios Thilikos visited the Paderborn group in February. M.J. Blesa attended the *Advanced Course on the Foundations of Computer Science* at the MPI, Saarbruecken, September 4-8, M.J. Blesa and J. Petit attended the *ALCOM-FT Summer School on Algorithm Engineering*, at the IASI-CNR, Rome, September 10-12.

In total, 24 ALCOM-FT reports appeared since the previous annual report [4, 9, 10, 13–17, 21–24, 57, 74, 75, 80–84, 107, 108, 135, 141]. Until now, a number of these have appeared in conference proceedings or journals, or are scheduled to do so soon. Among the journals are: Theoretical Computer Science, Random Structures, and Algorithms, Journal of Algorithms, ACM Computing Surveys. Researchers from the group have presented papers in: ICALP-2001, LICS-2002, MFCS-2001, SPAA-2001, WG-2002, ESA-2001, ALT-2001, IEEE Conf. on Data Mining-2001, NeuroCOLT-2002, COCOON-2001, COLT-2001, Discovery Science-2001, Formal Power Series-

2001, SAGA-2001, SIAM Workshop on Discrete Math and Data Mining-2001, Europar-2001, Edinburgh Workshop on Circuit and Proof Complexity-2001, GASCOM-2001.

Cologne

The activity at site Cologne is based in 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 payed full-time by ALCOM-FT funds.

During the second year of ALCOM-FT, Michael Jünger was Co-Chair of the organization and program committee of *Graph Drawing 2001* and member of the programm committee of the *Cologne Twente Workshop on Graphs and Combinatorial Optimization*.

Michael Jünger and Matthias Elf visited the MPI. Members of the site attended GD2001, OR2001, CP-AI-OR-2002 and the Aussois Workshop on Combinatorial Optimization. Michael Jünger gave an industrial talk at the Siemens AG (Zentralbereich Forschung und Entwicklung (ZFE)) in Munich. He was a teacher at the ALCOM-FT Summer School on Algorithm Engineering at the IASI-CNR in Rome. Matthias Elf also attended the school.

The ALCOM-FT reports [7,29,115,116] 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 and is complemented by a member from a new group (Philippe Robert, RAP) whose research focuses on networks and communication. The INRIA team of ALCOM-FT then comprises 5 permanent members (one junior and 4 senior members), 4 PhD students, and 4 associate members (from University of Caen and Paris). The research of two of the PhD students (M. Durand and V. Puyhaubert) takes place entirely within the framework of the ALCOM Project. The work packages concerned are WP1, WP2, WP4, and WP5. A postdoctoral position dedicated entirely to the cooperation within ALCOM-FT between INRIA and MPI has been set up for a full year (C. Banderier).

Research produced in the group continues to be 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 (WP4 and WP5), massive data sets (WP1), and large interconnection networks (WP2). In 2002, a synthesis report of some 200 pages has been produced [99]. We have developed an extensive visitor programme (with publication of seminar proceedings [63], 190 pages). New promising directions in the random generation of combinatorial structures are being explored [85]. Applications to bio-informatics and networks are being developed jointly with the Research Division of the French Telecom. We continue to play a major rôle in the organization of the *International Seminars on Analysis of Algorithms* and in the setting up of the French network ALÉA that is specifically dedicated to random structures and analysis of algorithms.

The INRIA site has produced 16 ALCOM-FT reports [25–28,39,63,85–88,99,106,113,114,118,156]. Of these, 7 have been accepted or are appearing in international journals (*Applied Prob.*, *Inf. Proc. Letters*, *Molecular Biology*, *Theoretical Comp. Sc.*), and 4 in international conferences (*ICALP*, *Trees*, *FPSAC*).

MPI

The working group Algorithms and Complexity at the Max-Planck-Institute for Computer science (MPII) participates in ALCOM-FT. MPII is involved in Work Packages 1–5. The Algorithms and Complexity group comprises its director, Prof. Kurt Mehlhorn, 17 researcher with Ph.D. (3

with habilitation), and 13 Ph.D. students. More than half of them work partially on ALCOM-FT. Ulrich Meyer, and from May 1, 2002, also Roman Dementiev, are paid from ALCOM-FT funds.

During the second year of ALCOM-FT, two Ph.D. theses related to the project were completed. Piotr Krysta worked on Approximation Algorithms for Combinatorial Optimization Problems in Graph Coloring and Network Design and Sven Thiel on narrowing Algorithms for Constraint Programming.

Several of our researchers and postdocs come from other ALCOM-FT sites; Berthold Vöcking from Paderborn, Dimitris Fotakis and Spyros Kontogiannis from CTI, and Cyril Banderier from INRIA. Visitors to the site include Artur Czumaj, Naveen Garg, Bruce Maggs, Peter Bro Miltersen (Aarhus), Paul Spirakis (CTI), Emo Welzl.

Members of the site have participated in most of the important relevant conferences, including ESA, FOCS, ICALP, SODA, SPAA, STOC, SWAT, and WAE. The site hosted the second Max-Planck Advanced Course on the Foundations of Computer Science.

The 30 ALCOM-FT reports [1, 8, 11, 20, 25, 27, 28, 30, 32, 33, 53, 66, 67, 69, 90, 96, 104, 117, 120, 122, 136–138, 145–150, 154] 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 second period, the site was involved in all five Work Packages of the ALCOM-FT project. In total, the two groups currently comprise 8 associate/assistant professors and 23 Ph.D. students. Ulf-Peter Schroeder and Martin Ziegler are paid part-time by ALCOM-FT funds.

During the second year of ALCOM-FT, Stephan Blazy completed his Ph.D. with the thesis *Numerical approximation of the Stokes equation using artificial boundary conditions in a system of pipes*. Visitors to the site include Artur Czumaj (University of Newark, USA), Birgitta Weber (University of Zurich, Switzerland), Christian Scheideler (University of Baltimore, USA), Dimitrios Thilikos (University of Barcelona, Spain), Grzgorz Malewicz (University of Boston, USA), Han Hoogeveen (University of Utrecht, Netherlands), Imrich Vrto (University of Bratislava, Slovakia), Jacques Bahi (University of Franche-Caute, France), Marios Mavronicolas (University of Cyprus, Cyprus), Martin Gairing (University of Clemson, USA), Paul Spirakis (University of Patras, Greece), Peter Ruzicka (University of Bratislava, Slovakia), Rastislav Kralovic (University of Bratislava, Slovakia).

Within the last year Burkhard Monien visited the ALCOM partners in Cyprus. Friedhelm Meyer auf der Heide visited the ALCOM partners in Rom. Torsten Fahle, Stefan Schamberger, Stephan Blazy, and Ulf-Peter Schroeder visited the ALCOM partners in Warwick. Klaus Volbert, Rolf Wanka, and Olaf Bonorden have participated at the ALCOM-FT-School on Algorithmic Engineering.

Moreover, within the last year Burkhard Monien has been a member of the program committees of EUROPAR'01, HiPC'01 (Program Chair), and EUROPAR'02 (Program Chair). Friedhelm Meyer auf der Heide was involved in the scientific program committees of ESA'01 (Program Chair), ARACNE'01, SIROCCO'02, EUROPAR'02, EUROMICRO'02, and PASA'02. Members of the site Paderborn have participated in most of the important and relevant conferences, including CP, CPAIOR, ESA, EUROPAR, IPDPS, MFCS, PDPTA, SIROCCO, SOFSEM, SPAA, STACS, and WG. Furthermore, several members of the site Paderborn have given talks on algorithms and our deliverables for industrial partners. Especially, the site has arranged several events within the forum *Industry meets Computer Science*, which is an initiative between the University of Paderborn and the society for business development in Paderborn.

The 28 ALCOM-FT reports [1, 34–36, 38, 43, 51, 52, 68, 91, 92, 97, 98, 100–103, 121, 131–133, 139, 143, 144, 152, 153, 156, 157] 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 WPs 1, 2 and 5. The group currently consists of 7 faculty members/postdoctoral researchers and 11 Ph.D. students.

During the second year of ALCOM-FT, the Ph.D. Thesis of I. Caragiannis ("Graph Colorings and their Applications to the Efficient Wavelength Routing in WDM Optical Networks") has been completed. One new student (S. Athanasopoulos) has begun his Ph.D.

Members of the site have participated in the Program Committees of major relevant Conferences: P. Spirakis in STOC 2002, COCOON 2002, RANDOM 2002, IFIP-TCS 2002, EUROPAR 2002, C. Zaroliagis in ICALP 2002, ESA 2002, WG 2002, S. Nikoletseas in EUROPAR 2002. C. Kaklamanis was Co-Chair of SIROCCO 2002. CTI organized SIROCCO 2002 in Andros, Greece.

P. Spirakis has visited MPI for 6 months by being awarded a position of "MPI distinguished visiting scientist". He was keynote speaker at the Perspectives in Computer Science (PICS 2002) Workshop, organized by Max-Planck Society in Berlin. Also, he was invited speaker at the British Colloquium on Theoretical Computer Science (BCTCS 2002) Meeting.

18 ALCOM-FT Technical reports [11,12,44,54–56,59–62,104,119,124–127,151,158] have been authored or co-authored by researchers of the site during the second year. Currently, 9 out of these (as well as one paper reported during the first year) have been accepted by conferences (including ICALP, MFCS, ACM Principles of Mobile Computing-POMC, WG, ALENEX) or journals.

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 consist of 8 faculty (full/associate/assistant professors) and 2 Ph.D. students. Camil Demetrescu is paid as researcher on ALCOM-FT funds; Paolo Terrevoli and Silvana Di Vincenzo are assisting with the coordination of the site.

During the second year of ALCOM-FT, Andrea Vitaletti has completed his Ph.D. with the thesis *Scheduling Algorithms and Localization Tools for Wireless Internet*. Currently, Luigi Laura (Rome "La Sapienza") and Fabrizio Grandoni (Rome "Tor Vergata") are working towards completing their Ph.D. programs. Fabrizio Grandoni is currently visiting the MPI site. Visitors to our site include Friedhelm Meyer auf der Heide, Devdatt Dubhashi, Kirk Pruhs, Stephen Alstrup, Adi Rosen, Franco P. Preparata, and Martin Tompa.

Members of the site have participated in most of the important relevant conferences, including ALENEX, FOCS, ICALP, SODA, SPAA, and STOC, and attended workshops at Dagstuhl and Oberwolfach. Several members of the group have been involved in program committees of various scientific events. Alberto Marchetti Spaccamela has been in the PC of WG'02, ESA'02 and TCS'02. Giuseppe F. Italiano has been member of the SWAT'02 PC, of the AIRO'02 PC, and of the ICTCS'01 PC. Giorgio Ausiello has been in the APPROX'02 PC, Stefano Leonardi in the ESA 01 PC, and Camil Demetrescu in the ESA'02 PC.

The site organizes permanent seminars on algorithms, one of which (SIA) in cooperation with the Department of Computer Science of the University of Roma "La Sapienza". See: `http://www.disp.uniroma2.it/algor/` and `http://www.dis.uniroma1.it/~algo`.

In September 2001, the site has hosted a Summer School on Algorithm Engineering, (see `http://www.info.uniroma2.it/~italiano/School`) and the first Project Workshop and Review Meeting. During 2002 the site is involved in the organization of ALGO 02, a federated event including the ESA, APPROX and WABI conferences. The site also organizes the School and Workshop on Models and Algorithms for the Web, which will take place in Udine in June 2002.

The 8 ALCOM-FT reports [18,19,31,58,76–79] have been authored or co-authored by members of the site. Until now, 7 of these have appeared in conference proceedings or journals, or are scheduled to do so soon.

Utrecht

The researchers participating in ALCOM-FT at the Utrecht site are members of the Algorithm Design group, which is part of the Institute of Information and Computing Sciences at Utrecht University. On January 1, 2002 the group was reconfigured into the 'Center for Algorithmic Systems', headed by Professors Jan van Leeuwen and Linda van der Gaag. The algorithms group is involved in WP 3 and 4. The group consists of 7 faculty members/researchers and 5 PhD students. In the second project year Dr K.M.J de Bontridder was appointed as a postdoc on the project and Dipl-Math Peter Lennartz joined the group as a PhD researcher (as of December 1, 2001). Koen de Bontridder will leave on July 1, 2002 to a job as algorithm specialist at Siemens VDO Dayton.

During the second year the research work on WP 3 and 4, in particular on tasks 3.1 and 4.3, got momentum because of the successful hiring of the postdoc and several PhD students. Nevertheless, part of the anticipated work in year one could only begin in the course of year two for this reason. Dr Han Hoogeveen visited ALCOM-FT partner Paderborn for one week in February 2002 for research cooperation. Peter Lennartz participated in the 'Spring School on Approximation Algorithms for Hard Problems' in Berlin. The group keeps close contacts with industry, especially with the National Aerospace Laboratory (NLR) and Philips Research.

Members of the group have participated actively in a number of international conferences and workshops, including ICALP, IPCO, MAPSP, SOFSEM and WG. In an excellent cooperation with ALCOM-FT sites Barcelona and Patras, Jan van Leeuwen (Utrecht) chaired the algorithm-track of the ICALP'2001 conference, which co-located with ACM-SIGACT's STOC'2001 conference in a unique event. Members of the group have authored or co-authored the ALCOM-FT reports [2, 3, 37, 71–73, 123]. Several of these reports have appeared in conference proceedings or journals or have been submitted.

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 four permanent academic staff (Leslie Goldberg, Paul Goldberg, Mike Paterson and Alex Tiskin), one research fellow (Russ Martin), and two Ph.D. students, Graham Cormode and Steven Kelk. Graham has recently completed the writing of his thesis and Steven will be finishing soon. For about half of the past year we had two research fellows, Petra Berenbrink (from Paderborn) and Tom Friedetzky. Friedetzky was paid by the ALCOM-FT project. Tom Friedetzky and Petra Berenbrink have now left to take up positions in Vancouver, Canada.

Visitors to Warwick included Micah Adler (U. Mass), Mary Cryan (Aarhus and Leeds), Artur Czumaj (New Jersy IT (ex Paderborn)), Martin Dyer (Leeds), Mark Jerrum (Edinburgh), Torsten Fahle (Paderborn), Uri Zwick (Tel Aviv) and Muthu Muthukrishnan (AT&T and Rutgers).

In May 2002 Warwick organized the *UKCRC Algorithms and Complexity Day* at which the main speakers were Claire Kenyon (Université Paris-Sud), Uri Zwick, and Muthu Muthukrishnan. In addition there were 5 contributed talks covering a broad range of topics.

Warwick ALCOMers have participated in many of the important relevant conferences, including SPAA 2001, STOC 2001, ICALP 2001, COLT 2001, Euro-Par 2001, FOCS 2001, STOC 2002, SODA 2002, British Combinatorial Conference 2001, ECOTEL 2001, and workshops NeuroCOLT 2002, RAND-APX 2002, and at least three Dagstuhl Seminars.

Cyprus

The activity at site Cyprus originates from the Theory group of the Department of Computer Science, University of Cyprus. The group is affiliated with the newly established Laboratory for Foundations of Computer Systems and Theoretical Computer Science. The group currently consists of 2 faculty members (M. Mavronicolas, A. Philippou), 1 associated faculty member (Y. Dimopoulos), 1 research scientist, 1 Ph.D. student and 2 Master students. The site is involved in WPs 2 and 4.

During the second year of ALCOM-FT, one new student (G. Hadjipollas) has begun his PhD. The site hosted short visits by three collaborating researchers: Burkhard Monien (University of Paderborn), Insup Lee (University of Pennsylvania, USA) and Costas Busch (Rensselaer Polytechnic Institute, USA).

Members of the site have participated in the Program Committees of major relevant Conferences: M. Mavronicolas in ICDCS 2002, Euro-Par 2002 (as Local Chair on Distributed Systems and Algorithms), and SIROCCO 2002.

M. Mavronicolas has been invited to deliver a 3-hour tutorial (titled "Game Theory in Network Routing: A Primer") in Euro-Par 2002. A. Philippou has co-delivered a 3-hour tutorial (titled "A Family of Resource-Bound Process Algebras for Modeling and Analysis of Embedded Systems") at ETAPS 2002.

7 ALCOM-FT Technical reports [104, 124, 125, 127, 130, 142, 142] have been authored or co-authored by researchers of the site during the second year. Currently, 5 out of these have been accepted for publication in conferences and journals (ICALP, SIROCCO, FORTE, IPL, and CCEJ).

Part III Information on Resources Used

In the following, we tabulate the number of person months spent at each site on each of the tasks of the ALCOM-FT project. The tables show the estimated (E) and actual (A) figures for year two alone, as well as the accumulated figures for year one and two. Only non-zero entries are shown.

Year Two

Task	Aarhus		Barcelona		Cologne		INRIA		MPI		Paderborn		CTI		Rome		Utrecht		Warwick		Cyprus	
	E	A	E	A	E	A	E	A	E	A	E	A	E	A	E	A	E	A	E	A	E	A
1.1									3	2	3	5			1	1			0.5	1		
1.2	3	4					0.5	1	2	6			3	3					2.5	2		
1.3			1	2			0.5	1		6					2	1			0.5	0.5		
1.4			3	3			0.5	1														
1.5																						
2.1													3	2								
2.2										3			4	3	1	2						
2.3													1	1								
2.4		2		1			1	3	4	6			2	1					4	3		
2.5							0.5	1	1	1	5	7			1	2			3	4	2	2
2.6									1	1					1	2						
2.7											3	5										
3.1	2	3									2	2					18	19				
3.2					5	5					1	3					2	3				
3.3											5	5					1	1				
4.1	3	4	5	5			3	18	5	15			3	2					6	9	1	1
4.2	3	4			12	12			2	3	1	3	1	1	2	6	3	2				
4.3	3	4	3	3			1	2	3	8	2	3	3	2	2	6	1	2	3	3		
5.1	3	3							3	3	4	4	2	2								
5.2							0.5	1	1				3	3								
5.3			5	5	3	3	1	3		3												
5.4				0.5					1	1					2	5						
5.5									3	3					2	2						
6.1									1	1	0.5	0.5	1	1	0.5	0.5		1				
6.2	3	4	0.5	0.5	1	1	0.5	1	0.5	0.5	0.5	0.5	1	1	0.5	0.5	0.5	0.5	1	1		
Sum	20	28	17.5	20	21	21	9	32	30.5	62.5	27	38	27	22	15	28	25.5	28.5	20.5	23.5	3	3

Note: The rather large actual figures compared to estimated figures for some sites are due to the fact that the estimated figures (which are in agreement with the figures in Annex 1 of the contract) correspond to full-time personnel only, whereas the actual figures include on-site Ph.D. students as well as researchers from other institutions that are associated to the site and accomplish part of their research there.

Accumulated for Year One and Two

Task	Aarhus		Barcelona		Cologne		INRIA		MPI		Paderborn		CTI		Rome		Utrecht		Warwick		Cyprus	
	E	A	E	A	E	A	E	A	E	A	E	A	E	A	E	A	E	A	E	A	E	A
1.1							0.5	1	6	7	6	7			2	2			1	2		
1.2	6	7					1	3	5	11			6	6					5	5		
1.3			3	4			0.5	1		6					4	2			1	0.5		
1.4			6	6			0.5	1		0.5												
1.5									2	4												
2.1													6	6	1	1						
2.2										3			8	9	2.5	5						
2.3									3	5			2	3								
2.4		2		1			2	7	4	6			4	4					9	6		
2.5							1	3	3	3	10	12			2.5	5			5	7	2	2
2.6									3	3	4	3.5			1	2						
2.7											6	8										
3.1	2	3									2	2					27	21				
3.2					7	7					1	3					5	4				
3.3											10	10					2	1.5				
4.1	6	8	10	11			6	32	9	20			6	7					12	17	1	1
4.2	7	8			24	24			5	7	2	4	2	2	4	12	5	4				
4.3	6	8	6	7			2	4	6	12	4.5	5	6	6	3	9	2	3	6.5	7		
5.1	6	6							5	5.5	6	5	4	4	1	1						
5.2							1	2	2	0.5			6	6								
5.3			10	10	6	6	2	5	1	4												
5.4				0.5					1	1					4	9						
5.5									4	4					3	5						
6.1	1	1							3	3	0.5	1.5	2	1	1	1		1				
6.2	6	8	2.5	1.5	2	2	1.5	2	2	2	1.5	1.5	2	2	1	1	1	1	2	2		
Sum	40	51	37.5	41	39	39	18	61	64	107.5	53.5	62.5	54	56	30	55	42	35.5	41.5	46.5	3	3

Part IV Scientific Reports

In the following, we list all scientific reports which have been published in the ALCOM-FT Technical Report Series in the period June 1, 2001 to May 31, 2002. All reports are available online at `www.brics.dk/ALCOM-FT/TR`.

References

[1] Micah Adler, Harald Räcke, Naveen Sivadasan, Christian Sohler, and Berthold Vöcking. Randomized Pursuit-Evasion in Graphs. Technical Report ALCOMFT-TR-02-35, ALCOM-FT, Paderborn and MPI, May 2002. Appearance: In *"Proceedings of the 29th International Colloquium on Automata, Languages and Programming (ICALP)"*, 2002.

[2] Marjan van den Akker and Han Hoogeveen. Minimizing the number of late jobs in case of stochastic processing times. Technical Report ALCOMFT-TR-02-83, ALCOM-FT, Utrecht, May 2002.

[3] Marjan van den Akker, Han Hoogeveen, and Gerhard J. Woeginger. The two-machine open shop problem: to fit, or not to fit, that was the question. Technical Report ALCOMFT-TR-02-82, ALCOM-FT, Utrecht, May 2002.

[4] E. Alba, F. Almeida, M. Blesa, J. Cabeza, C. Cotta, M. Díaz, I. Dorta, J. Gabarró, C. León, J. Luna, L. Moreno, C. Pablos, J. Petit, A. Rojas, and F. Xhafa. MALLBA: A library of skeletons for combinatorial optimisation. Technical Report ALCOMFT-TR-02-120, ALCOM-FT, Barcelona, May 2002. Appearance: In *Euro-Par 2002*. Springer-Verlag, LNCS, 2002. To appear.

[5] Susanne Albers, Lene M. Favrholdt, and Oliver Giel. On Paging with Locality of Reference. Technical Report ALCOMFT-TR-02-56, ALCOM-FT, Aarhus, May 2002. Appearance: In *34th Annual ACM Symposium on the Theory of Computing*, 2002. To appear.

[6] Stephen Alstrup, Gerth Stølting Brodal, Inge Li Gørtz, and Theis Rauhe. Time and Space Efficient Multi-Method Dispatching. Technical Report ALCOMFT-TR-02-76, ALCOM-FT, Aarhus, May 2002. Appearance: In *Proc. 8th Scandinavian Workshop on Algorithm Theory*, Lecture Notes in Computer Science. Springer-Verlag, Berlin, 2002.

[7] Ernst Althaus, Alexander Bockmayr, Matthias Elf, Michael Jünger, Thomas Kasper, and Kurt Mehlhorn. SCIL - Symbolic Constraints in Integer Linear Programming. Technical Report ALCOMFT-TR-02-133, ALCOM-FT, Cologne and MPI, May 2002.

[8] Ernst Althaus and Christian Fink. A Polyhedral Approach to Surface Reconstruction from Planar Contours. Technical Report ALCOMFT-TR-02-88, ALCOM-FT, MPI, May 2002.

[9] Carme Àlvarez, Maria Blesa, and Maria Serna. A characterization of universal stability for directed graphs in the adversarial queueing model. Technical Report ALCOMFT-TR-02-104, ALCOM-FT, Barcelona, May 2002.

[10] Carme Àlvarez, Maria Blesa, and Maria Serna. Universal stability of undirected graphs in the adversarial queueing model. Technical Report ALCOMFT-TR-02-103, ALCOM-FT, Barcelona, May 2002. Appearance: In *Proceedings of the fourteenth annual ACM Symposium on Parallel algorithms and architectures, SPAA 2002*, 2002. to appear.

[11] Maria Andreou, Dimitris Fotakis, Sotiris Nikoletseas, Vicky Papadopoulou, and Paul Spirakis. On Radiocoloring Hierarchically Specified Planar Graphs: PSPACE-completeness and Approximations. Technical Report ALCOMFT-TR-02-68, ALCOM-FT, Patras and MPI,

May 2002. Appearance: In *27th International Symposium on Mathematical Foundations of Computer Science, MFCS 2002*. Springer-Verlag, 2002.

[12] Maria Andreou and Paul Spirakis. The Square of Any Planar Graph is an Almost Perfect Graph. Technical Report ALCOMFT-TR-02-89, ALCOM-FT, Patras, May 2002. Appearance: Under review.

[13] Albert Atserias. Improved Bounds for the Weak Pigeonhole Principle and Infinitely Many Primes from Weaker Axioms. Technical Report ALCOMFT-TR-02-122, ALCOM-FT, Barcelona, May 2002. Appearance: *Theoretical Computer Science*, 2002. To appear.

[14] Albert Atserias. The Complexity of Resource-Bounded Propositional Proofs. Technical Report ALCOMFT-TR-02-132, ALCOM-FT, Barcelona, May 2002.

[15] Albert Atserias. Unsatisfiable Random Formulas are Hard to Certify. Technical Report ALCOMFT-TR-02-129, ALCOM-FT, Barcelona, May 2002. Appearance: In *17th IEEE Symposium on Logic in Computer Science*, 2002. To appear.

[16] Albert Atserias, Maria Luisa Bonet, and Juan Luis Esteban. Lower Bounds for the Weak Pigeonhole Principle and Random Formulas Beyond Resolution. Technical Report ALCOMFT-TR-02-23, ALCOM-FT, Barcelona, May 2002. Appearance: *Information and Computation*, 2002. To appear.

[17] Albert Atserias, Nicola Galesi, and Pavel Pudlák. Monotone Simulations of Nonmonotone Proofs. Technical Report ALCOMFT-TR-02-131, ALCOM-FT, Barcelona, May 2002. Appearance: *Journal of Computer and System Sciences*, 2002. To appear.

[18] Giorgio Ausiello, Paolo G. Franciosa, and Daniele Frigioni. Directed Hypergraphs: Problems, Algorithmic Results, and a Novel Decremental Approach. Technical Report ALCOMFT-TR-02-5, ALCOM-FT, Rome, March 2002. Appearance: In *ICTCS: Italian Conference on Theoretical Computer Science*, volume 2202 of *Lecture Notes in Computer Science*, pages 312–327. Springer, 2001.

[19] Giorgio Ausiello, Paolo G. Franciosa, and Daniele Frigioni. Partially Dynamic Maintenance of Minimum Weight Hyperpaths. Technical Report ALCOMFT-TR-02-11, ALCOM-FT, Rome, April 2002.

[20] David A. Bader, Bernard M.E. Moret, and Peter Sanders. High-Performance Algorithm Engineering for Parallel Computation. Technical Report ALCOMFT-TR-02-29, ALCOM-FT, MPI, May 2002.

[21] R. Baeza-Yates, J. Gabarró, and X. Messeguer. Fringe Analysis of Synchronized Parallel Insertion Algorithms in 2–3 Trees. Technical Report ALCOMFT-TR-02-119, ALCOM-FT, Barcelona, May 2002. Appearance: *Theoretical Computer Science*, 2002. To appear.

[22] J Baixeries and G Casas-Garriga. Sampling Strategies for Finding Frequent Sets. Technical Report ALCOMFT-TR-02-6, ALCOM-FT, Barcelona, April 2002.

[23] José Balcázar. Rules with Bounded Negations and the Coverage Inference Scheme. Technical Report ALCOMFT-TR-02-115, ALCOM-FT, Barcelona, May 2002.

[24] José L. Balcázar, Yang Dai, and Osamu Watanabe. Provably fast support vector regression using random sampling. Technical Report ALCOMFT-TR-02-84, ALCOM-FT, Barcelona, May 2002. Appearance: In *Workshop Discrete Mathematics and Data Mining at SIAM DM Conference*, Arlington, USA, april 2002, 2002.

[25] Cyril Banderier. Limit laws for basic parameters of lattice paths with unbounded jumps. Technical Report ALCOMFT-TR-02-22, ALCOM-FT, INRIA and MPI, May 2002. Appearance: In *Trends in Mathematics (Versailles, 2002)*, Basel, 2002. Birkhäuser. To appear.

[26] Cyril Banderier and Philippe Flajolet. Basic Analytic Combinatorics of Directed Lattice Paths. Technical Report ALCOMFT-TR-01-188, ALCOM-FT, INRIA, December 2001.

[27] Cyril Banderier, Kurt Mehlhorn, and Rene Beier. Average case analysis of three combinatorial algorithms under limited randomness. Technical Report ALCOMFT-TR-02-21, ALCOM-FT, INRIA and MPI, May 2002. Appearance: In *submitted to FOCS'02*, 2002.

[28] Cyril Banderier and Donatella Merlini. Lattice paths with an infinite set of jumps. Technical Report ALCOMFT-TR-02-24, ALCOM-FT, INRIA and MPI, May 2002. Appearance: In *Proceedings of Formal Power Series and Algebraic Combinatorics'02*, 2002.

[29] Wilhelm Barth, Michael Jünger, and Petra Mutzel. Simple and Efficient Bilayer Cross Counting. Technical Report ALCOMFT-TR-02-125, ALCOM-FT, Cologne, May 2002.

[30] Hannah Bast. Scheduling at Twilight the Easy Way. Technical Report ALCOMFT-TR-02-36, ALCOM-FT, MPI, May 2002.

[31] Luca Becchetti, Suhas Diggavi, Stefano Leonardi, Alberto Marchetti-Spaccamela, Muthu Muthukrishnan, Thiaga Nandagopal, and Andrea Vitaletti. Parallel scheduling problems in next generation wireless networks. Technical Report ALCOMFT-TR-02-93, ALCOM-FT, Rome, May 2002.

[32] René Beier, Peter Sanders, and Naveen Sivadasan. Energy Optimal Routing in Radio Networks Using Geometric Data Structures. Technical Report ALCOMFT-TR-02-32, ALCOM-FT, MPI, May 2002. Appearance: In *29th International Colloquium on Automata, Languages and Programming*. Springer, 2002. to appear.

[33] Nicolas Beldiceanu, Qi Guo, and Sven Thiel. Non-overlapping Constraints between Convex Polytopes. Technical Report ALCOMFT-TR-02-47, ALCOM-FT, MPI, May 2002. Appearance: In Toby Walsh, editor, *Proceedings of the 7th International Symposium on Principles and Practice of Constraint Programming (CP-01)*, volume 2239 of *Lecture Notes in Computer Science*, pages 392–407, Paphos, Cyprus, November 2001. Springer.

[34] Sergei Bezrukov and Robert Elsässer. Edge-Isoperimetric Problems for Cartesian Powers of Regular Graphs. Technical Report ALCOMFT-TR-02-97, ALCOM-FT, Paderborn, May 2002. Appearance: In *Proceedings of 27th International Workshop on Graph-Theoretic Concepts in Computer Science (WG)*, volume 2204 of *LNCS*, pages 9–20. Springer, 2001.

[35] Hans L. Bodlaender and Fedor V. Fomin. Approximation of pathwidth of outerplanar graphs. Technical Report ALCOMFT-TR-02-12, ALCOM-FT, Paderborn and Utrecht, May 2002. Appearance: *J. Algorithms*, page to appear, 2002.

[36] Hans L. Bodlaender and Fedor V. Fomin. Tree decompositions with small cost. Technical Report ALCOMFT-TR-02-14, ALCOM-FT, Paderborn and Utrecht, May 2002. Appearance: In *Proceedings of the 8th Scandinavian Workshop on Algorithm Theory (SWAT 2002)*, page to appear. Springer, 2002.

[37] Hans L. Bodlaender and Udi Rotics. Computing the treewidth and the minimum fill-in with the modular decomposition. Technical Report ALCOMFT-TR-02-62, ALCOM-FT, Utrecht, May 2002.

[38] Olaf Bonorden, Friedhelm Meyer auf der Heide, and Rolf Wanka. Composition of Efficient NestedBSP Algorithms: Minimum Spanning Tree Computation as an InstructiveExample. Technical Report ALCOMFT-TR-02-46, ALCOM-FT, Paderborn, May 2002. Appearance: In *Proc. Int. Conf. on Parallel and Distributed Processing Techniques and Applications (PDPTA)*, 2002. to appear.

[39] Jérémie Bourdon and Brigitte Vallée. Generalized Pattern Matching Statistics. Technical Report ALCOMFT-TR-02-130, ALCOM-FT, INRIA, May 2002.

[40] Joan Boyar, Lene M. Favrholdt, Kim S. Larsen, and Morten N. Nielsen. Extending the Accommodating Function. Technical Report ALCOMFT-TR-02-4, ALCOM-FT, Aarhus, March 2002. Appearance: In *Proceedings of COCOON 2002*, 2002. To appear.

[41] Joan Boyar, Susan Krarup, and Morten N. Nielsen. Seat Reservation Allowing Seat Changes. Technical Report ALCOMFT-TR-02-3, ALCOM-FT, Aarhus, March 2002. Appearance: In *Proceedings of the Satellite Workshops of the 28th International Colloquium on Automata, Languages, and Programming*, pages 26–40, 2001.

[42] Joan Boyar, René Peralta, and Denis Pochuev. Concrete Conjunctive Complexity of Symmetric Functions. Technical Report ALCOMFT-TR-02-2, ALCOM-FT, Aarhus, January 2002.

[43] André Brinkmann, Kay Salzwedel, and Christian Scheideler. Compact, Adaptive Placement Schemes for Non-Uniform Distribution Requirements. Technical Report ALCOMFT-TR-02-124, ALCOM-FT, Paderborn, May 2002. Appearance: In *Proc. of the 14th ACM Annual Symp. on Parallel Algorithms and Architectures (SPAA'02)*, August 2002. To appear.

[44] Gerth Stølting Brodal, George Lagogiannis, Christos Makris, Athanasios Tsakalidis, and Kostas Tsichlas. Optimal Finger Search Trees in the Pointer Machine. Technical Report ALCOMFT-TR-02-77, ALCOM-FT, Aarhus and Patras, May 2002. Appearance: In *Proc. 34th Annual ACM Symposium on Theory of Computing*, 2002.

[45] Gerth Stølting Brodal, Rune B. Lyngsø, Anna Östlin, and Christian N. S. Pedersen. Solving the String Statistics Problem in Time $O(n \log n)$. Technical Report ALCOMFT-TR-02-55, ALCOM-FT, Aarhus, May 2002. Appearance: In *Proc. 29th International Colloquium on Automata, Languages, and Programming*, Lecture Notes in Computer Science. Springer Verlag, 2002. To appear.

[46] Gerth Stølting Brodal and Rolf Fagerberg. Cache Oblivious Distribution Sweeping. Technical Report ALCOMFT-TR-02-52, ALCOM-FT, Aarhus, May 2002. Appearance: In *Proceedings of ICALP 02*, 2002. To appear.

[47] Gerth Stølting Brodal and Rolf Fagerberg. Funnel Heap - A Cache Oblivious Priority Queue. Technical Report ALCOMFT-TR-02-136, ALCOM-FT, Aarhus, June 2002.

[48] Gerth Stølting Brodal, Rolf Fagerberg, and Riko Jacob. Cache Oblivious Search Trees via Binary Trees of Small Height. Technical Report ALCOMFT-TR-02-53, ALCOM-FT, Aarhus, May 2002. Appearance: In *Proc. 13th Annual ACM-SIAM Symposium on Discrete Algorithms,*, pages 39–48, 2002.

[49] Gerth Stølting Brodal, Rolf Fagerberg, and Christian N. S. Pedersen. Computing the Quartet Distance between Evolutionary Trees in Time $O(n \log n)$. Technical Report ALCOMFT-TR-02-54, ALCOM-FT, Aarhus, May 2002.

[50] Gerth Stølting Brodal and Riko Jacob. Time-dependent networks as models to achieve fast exact time-table queries. Technical Report ALCOMFT-TR-01-176, ALCOM-FT, Aarhus, September 2001.

[51] Hajo Broersma, Fedor. V Fomin, Jan Kratochvíl, and Gerhard J. Woeginger. Planar graph coloring with forbidden subgraphs: Why trees and paths are dangerous. Technical Report ALCOMFT-TR-02-15, ALCOM-FT, Paderborn, May 2002. Appearance: In *Proceedings of the 8th Scandinavian Workshop on Algorithm Theory (SWAT 2002)*, page to appear. Springer, 2002.

[52] Hajo Broersma, Fedor V. Fomin, Jaroslav Neˇ setřil, and Gerhard J. Woeginger. More about subcolorings. Technical Report ALCOMFT-TR-02-16, ALCOM-FT, Paderborn, May 2002. Appearance: In *Proceedings of the 28th International Workshop on Graph-Theoretic Concepts in Computer Science (WG 2002)*, page to appear. Springer, 2002.

[53] Stefan Burkhardt and Juha Kärkkäinen. One-gapped q-Gram Filters for Levenshtein Distance. Technical Report ALCOMFT-TR-02-41, ALCOM-FT, MPI, May 2002. Appearance: In *Proc. 13th Annual Symposium on Combinatorial Pattern Matching*. Springer, 2002. To appear.

[54] Ioannis Caragiannis, Afonso Ferreira, Christos Kaklamanis, Stephane Perennes, Pino Persiano, and Herve Rivano. Approximate Constrained Bipartite Edge Coloring. Technical Report ALCOMFT-TR-01-184, ALCOM-FT, Patras, October 2001.

[55] Ioannis Caragiannis, Christos Kaklamanis, and Panagiotis Kanellopoulos. New Bounds on the Size of the Minimum Feedback Vertex Set in Meshes and Butterflies. Technical Report ALCOMFT-TR-01-182, ALCOM-FT, Patras, October 2001.

[56] Ioannis Caragiannis, Christos Kaklamanis, and Evi Papaioannou. Randomized Call Control in Sparse Wireless Cellular Networks. Technical Report ALCOMFT-TR-01-183, ALCOM-FT, Patras, October 2001.

[57] G. Casas-Garriga. A Multi-Step Pruning Strategy to Remove Uninteresting Association Rules. Technical Report ALCOMFT-TR-02-121, ALCOM-FT, Barcelona, May 2002.

[58] Giuseppe Cattaneo, Pompeo Faruolo, Umberto Ferraro-Petrillo, and Giuseppe F. Italiano. Maintaining Dynamic Minimum Spanning Trees: An Experimental Study. Technical Report ALCOMFT-TR-02-99, ALCOM-FT, Rome, May 2002. Appearance: In *Proc. Workshop on Algorithm Engineering and Experimentaion*, 2002.

[59] Ioannis Chatzigiannakis and Sotiris Nikoletseas. An Adaptive Compulsory Protocol for Basic Communication in Highly Changing Ad hoc Mobile Networks. Technical Report ALCOMFT-TR-02-73, ALCOM-FT, Patras, May 2002. Appearance: In *Proceedings of 16th International Parallel & Distributed Processing Symposium, 2nd International Workshop on Parallel and Distributed Computing Issues in Wireless networks and Mobile Computing – IPDPS 2002 Workshops*. IEEE/ACM, 2002.

[60] Ioannis Chatzigiannakis, Sotiris Nikoletseas, and Paul Spirakis. On the Average and Worst-case Efficiency of Some New Distributed Algorithms for Communication and Control in Ad-hoc Mobile Networks. Technical Report ALCOMFT-TR-02-72, ALCOM-FT, Patras, May 2002. Appearance: In *Proceedings of the 1st Workshop on Principles of Mobile Computing – POMC 2001*. ACM, 2001. Invited Paper.

[61] Ioannis Chatzigiannakis, Sotiris Nikoletseas, and Paul Spirakis. Smart Dust Local Detection and Propagation Protocols. Technical Report ALCOMFT-TR-02-114, ALCOM-FT, Patras, May 2002.

[62] Ioannis Chatzigiannakis, Sotiris Nikoletseas, Paul Spirakis, and Christos Zaroliagis. Experimenting with Relay Protocols for Communication in Ad-hoc Mobile Networks. Technical Report ALCOMFT-TR-02-71, ALCOM-FT, Patras, May 2002.

[63] Fréderic Chyzak. Algorithms Seminar, 2000-2001. Technical Report ALCOMFT-TR-02-7, ALCOM-FT, INRIA, April 2002.

[64] M Cryan, M Dyer, L A Goldberg, M Jerrum, and R Martin. Rapidly Mixing Markov Chains for Sampling Contingency Tables with a Constant Number of Rows. Technical Report ALCOMFT-TR-02-19, ALCOM-FT, Warwick, May 2002.

[65] Mary Cryan and Peter Bro Miltersen. On Pseudorandom generators in NC^0. Technical Report ALCOMFT-TR-02-67, ALCOM-FT, Aarhus, May 2002. Appearance: In *Proc. 26th International Symposium on Mathematical Foundations of Computer Science*, volume 2136 of *Lecture Notes in Computer Science*, pages 272–284. Springer, 2001.

[66] Bela Csaba, Marek Karpinski, and Piotr Krysta. Approximability of Dense and Sparse Instances of Minimum 2-Connectivity, TSP and Path Problems. Technical Report ALCOMFT-TR-02-26, ALCOM-FT, MPI, May 2002. Appearance: In *Proc. 13th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 74–83, 2002.

[67] Artur Czumaj, Piotr Krysta, and Berthold Voecking. Selfish Traffic Allocation for Server Farms. Technical Report ALCOMFT-TR-02-25, ALCOM-FT, MPI, May 2002. Appearance: In *Proc. 34th ACM Symposium on Theory of Computing (STOC)*, 2002. To appear.

[68] Artur Czumaj and Christian Sohler. Property Testing with Geometric Queries. Technical Report ALCOMFT-TR-02-44, ALCOM-FT, Paderborn, May 2002. Appearance: In *Proceedings of the9th Annual European Symposium on Algorithms (ESA '01)*, pages 266–277, 2001.

[69] Artur Czumaj and Berthold Vöcking. Tight Bounds for Worst-Case Equilibria. Technical Report ALCOMFT-TR-02-27, ALCOM-FT, MPI, May 2002. Appearance: In *Proceedings of the 13th ACM Symp. on Discrete Algorithms (SODA)*, pages 413–420, 2002.

[70] Ivan Bjerre Damgård and Gudmund Skovbjerg Frandsen. An Extended Quadratic Frobenius Primality Test with Average Case Error Estimates. Technical Report ALCOMFT-TR-01-190, ALCOM-FT, Aarhus, December 2001.

[71] K.M.J. De Bontridder. A purchase lot sizing problem with quantity discounts. Technical Report ALCOMFT-TR-02-78, ALCOM-FT, Utrecht, May 2002.

[72] K.M.J. De Bontridder, B.V. Halldórsson, M.M. Halldórsson, J.K. Lenstra, R. Ravi, and l. Stougie. Approximation algorithms for the minimum test set problem. Technical Report ALCOMFT-TR-02-80, ALCOM-FT, Utrecht, May 2002.

[73] K.M.J. De Bontridder, J.K. Lenstra, J.B. Orlin, and L. Stougie. Branch and bound algorithms for the minimum test set problem. Technical Report ALCOMFT-TR-02-79, ALCOM-FT, Utrecht, May 2002.

[74] Erik D. Demaine, MohammadTaghi Hajiaghayi, and Dimitrios Thilikos. 1.5-Approximation for Treewidth of Graphs Excluding a Graph with One Crossing as a Minor. Technical Report ALCOMFT-TR-02-113, ALCOM-FT, Barcelona, May 2002.

[75] Erik D. Demaine, MohammadTaghi Hajiaghayi, and Dimitrios M. Thilikos. Exponential Speedup of Fixed Parameter Algorithms on $K_{3,3}$-minor-free or K_5-minor-free Graphs. Technical Report ALCOMFT-TR-02-112, ALCOM-FT, Barcelona, May 2002.

[76] C. Demetrescu, I. Finocchi, and J. T. Stasko. Specifying Algorithm Visualizations: Interesting Events or State Mapping? Technical Report ALCOMFT-TR-01-177, ALCOM-FT, Rome, October 2001.

[77] C. Demetrescu and G.F. Italiano. A New Approach to Dynamic All Pairs Shortest Paths. Technical Report ALCOMFT-TR-02-92, ALCOM-FT, Rome, May 2002.

[78] C. Demetrescu and G.F. Italiano. Improved Bounds and New Trade-Offs for Dynamic All Pairs Shortest Paths. Technical Report ALCOMFT-TR-02-1, ALCOM-FT, Rome, January 2002. Appearance: In *Proc. of the 29-th International Colloquium on Automata, Languages, and Programming (ICALP'02), Malaga, Spain*, 2002.

[79] Camil Demetrescu and Mikkel Thorup. Oracles for Distances Avoiding a Link-failure. Technical Report ALCOMFT-TR-01-175, ALCOM-FT, Rome, September 2001. Appearance: In *Proc. of the 13th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'02), San Francisco, California*, 2002.

[80] J. Díaz, N. Do, M. J. Serna, and N. C. Wormald. Bisection of Random Cubic Graphs. Technical Report ALCOMFT-TR-02-102, ALCOM-FT, Barcelona, May 2002.

[81] J. Díaz, J. Neˇ setřil, M. Serna, and D.M. Thilikos. H-colorings of Large Degree Graphs. Technical Report ALCOMFT-TR-02-117, ALCOM-FT, Barcelona, May 2002.

[82] Josep Díaz, Maria Serna, and Dimitrios Thilikos. The complexity of restrictive H-coloring. Technical Report ALCOMFT-TR-02-101, ALCOM-FT, Barcelona, May 2002. Appearance: In *Graph Theoretic Concepts in Computer Science, WG 2002*, Lecture NOtes in Computer Science. Springer, 2002. To appear.

[83] Josep Díaz, Maria Serna, and Dimitrios M. Thilikos. (H, C, K)-coloring: Fast, Easy and Hard cases. Technical Report ALCOMFT-TR-02-105, ALCOM-FT, Barcelona, May 2002. Appearance: In Jiri Sgall, Ales Pultr, and Petr Kolman, editors, *Mathematical Foundations of Computer Science 2001, 26th International Symposium, MFCS 2001 Marianske Lazne, Czech Republic, August 27-31, 2001, Proceedings*, volume 2136 of *Lecture Notes in Computer Science*, pages 304–315. Springer, 2001.

[84] Amalia Duch and Conrado Martínez. On the Average Performance of Orthogonal Range Search in Multidimensional Data Structures. Technical Report ALCOMFT-TR-01-185, ALCOM-FT, Barcelona, November 2001. Appearance: In *Proc. 29th Int. Col. Automata, Languages and Programming (ICALP)*, Lect. Notes in Comp. Sci. Springer, 2002. to appear.

[85] Philippe Duchon, Philippe Flajolet, Guy Louchard, and Gilles Schaeffer. Random Sampling from Boltzmann Principles. Technical Report ALCOMFT-TR-01-189, ALCOM-FT, INRIA, December 2001.

[86] Vincent Dumas, Fabrice Guillemin, and Philippe Robert. A Markovian analysis of AIMD algorithms. Technical Report ALCOMFT-TR-01-179, ALCOM-FT, INRIA, October 2001. Appearance: *Advances in Applied Probability*, 34(1):–, 2002.

[87] Vincent Dumas, Fabrice Guillemin, and Philippe Robert. Effective bandwidths in a multiclass priority queueing system. Technical Report ALCOMFT-TR-01-178, ALCOM-FT, INRIA, October 2001.

[88] Marianne Durand. Asymptotic Analysis of an Optimized Quicksort Algorithm. Technical Report ALCOMFT-TR-02-128, ALCOM-FT, INRIA, May 2002.

[89] M Dyer, L A Goldberg, and M Jerrum. Counting and sampling H-colourings. Technical Report ALCOMFT-TR-02-20, ALCOM-FT, Warwick, May 2002.

[90] Stefan Edelkamp and Ulrich Meyer. Theory and Practice of Time-Space Trade-Offs in Memory Limited Search. Technical Report ALCOMFT-TR-02-51, ALCOM-FT, MPI, May 2002. Appearance: In *Proc. Joint German/Austrian Conference on Artificial Intelligence (KI-2001)*, volume 2174 of *LNAI*, pages 169–184. Springer, 2001.

[91] Robert Elsässer, Burkhard Monien, and Robert Preis. Diffusion Schemes for Load Balancing on Heterogeneous Networks. Technical Report ALCOMFT-TR-02-96, ALCOM-FT, Paderborn, May 2002. Appearance: *Theory of Computing Systems*, page to appear, 2002.

[92] Robert Elsässer, Burkhard Monien, Stefan Schamberger, and Günter Rote. Toward optimal diffusion matrices. Technical Report ALCOMFT-TR-02-98, ALCOM-FT, Paderborn, May 2002. Appearance: In A. Ferreira and A. Gottlieb, editors, *Proceedings of 16th International Parallel and Distributed Processing Symposium (IPDPS)*, pages 67 plus CD–ROM. IEEE Computer Society, 2002.

[93] Leah Epstein and Lene M. Favrholdt. On-Line Maximizing the Number of Items Packed in Variable-Sized Bins. Technical Report ALCOMFT-TR-02-59, ALCOM-FT, Aarhus, May 2002. APPEARANCE: In *Eighth Annual International Computing and Combinatorics Conference*, 2002. To appear.

[94] Leah Epstein and Lene M. Favrholdt. Optimal Non-Preemptive Semi-Online Scheduling on Two Related Machines. Technical Report ALCOMFT-TR-02-58, ALCOM-FT, Aarhus, May 2002. APPEARANCE: In *27th International Symposium on Mathematical Foundations of Computer Science*, 2002. To appear.

[95] Leah Epstein and Lene M. Favrholdt. Optimal Preemptive Semi-Online Scheduling to Minimize Makespan on Two Related Machines. Technical Report ALCOMFT-TR-02-57, ALCOM-FT, Aarhus, May 2002. APPEARANCE: *Operations Research Letters*, 2002. To appear.

[96] Leah Epstein, Csanád Imreh, and Rob van Stee. More on weighted servers or FIFO is better than LRU. Technical Report ALCOMFT-TR-02-28, ALCOM-FT, MPI, May 2002.

[97] Torsten Fahle. Cost Based Filtering vs. Upper Bounds for Maximum Clique. Technical Report ALCOMFT-TR-02-106, ALCOM-FT, Paderborn, May 2002. APPEARANCE: In *Proceedings ESA'02, to appear*, 2002. improved version of a paper published in the Proceedings of the CP-AI-OR 2002.

[98] Torsten Fahle, Stefan Schamberger, and Meinolf Sellmann. Symmetry Breaking. Technical Report ALCOMFT-TR-02-91, ALCOM-FT, Paderborn, May 2002. APPEARANCE: In *7th Internat. Conference on the Principles and Practice of Constraint Programming – CP'01*, volume 2239 of *LNCS*, pages 93–107. Springer, 2001.

[99] Philippe Flajolet and Robert Sedgewick. Analytic Combinatorics, Symbolic Combinatorics. Technical Report ALCOMFT-TR-02-123, ALCOM-FT, INRIA, May 2002.

[100] Fedor V. Fomin, Dieter Kratsch, and Haiko Müller. Algorithms for graphs with small octopus. Technical Report ALCOMFT-TR-01-181, ALCOM-FT, Paderborn, October 2001.

[101] Fedor V. Fomin, Dieter Kratsch, and Jean-Christophe Novelli. Approximating minimum cocolourings. Technical Report ALCOMFT-TR-01-180, ALCOM-FT, Paderborn, October 2001. APPEARANCE: In *Proceedings of the 13th International Symposium on Fundamentals of Computation Theory (FCT)*, pages 118–125, 2001.

[102] Fedor V. Fomin and Andrzej Lingas. Approximation algorithms for time-dependent orienteering. Technical Report ALCOMFT-TR-02-13, ALCOM-FT, Paderborn, May 2002. APPEARANCE: *Inform. Proc. Letters*, page to appear, 2002.

[103] Fedor V. Fomin, Martín Matamala, and Ivan Rapaport. The complexity of approximating the oriented diameter of chordal graphs. Technical Report ALCOMFT-TR-02-17, ALCOM-FT, Paderborn, May 2002. APPEARANCE: In *Proceedings of the 28th International Workshop on Graph-Theoretic Concepts in Computer Science (WG 2002)*, page to appear. Springer, 2002.

[104] Dimitris Fotakis, Spyros Kontogiannis, Elias Koutsoupias, Marios Mavronicolas, and Paul Spirakis. The Structure and Complexity of Nash Equilibria for a Selfish Routing Game. Technical Report ALCOMFT-TR-02-37, ALCOM-FT, Patras and MPI and Cyprus, May 2002. APPEARANCE: In *29th International Colloquium on Automata, Languages and Programming (ICALP'02)*, page TO APPEAR, 2002.

[105] Jens S. Frederiksen and Kim S. Larsen. Packet Bundling. Technical Report ALCOMFT-TR-02-60, ALCOM-FT, Aarhus, May 2002. APPEARANCE: In *Eighth Scandinavian Workshop on Algorithm Theory*, 2002. To appear.

[106] Christine Fricker, Philippe Robert, and Danielle Tibi. A degenerate central limit theorem for single resource loss systems. Technical Report ALCOMFT-TR-02-9, ALCOM-FT, INRIA, April 2002.

[107] J. Gabarró, A. Stewart, and M. Clint. Grab and Go Systems: a CPO approach to concurrent web and grid-based computation. Technical Report ALCOMFT-TR-02-118, ALCOM-FT, Barcelona, May 2002.

[108] Ricard Gavaldà and Osamu Watanabe. Sequential Sampling Algorithms: Unified Analysis and Lower Bounds. Technical Report ALCOMFT-TR-02-116, ALCOM-FT, Barcelona, May 2002. Appearance: In *Proc. 1st Intl. Symposium on Stochastic Algorithms: Foundations and Applications (SAGA'01).*, volume 2264. Springer-Verlag Lecture Notes in Computer Science, 2001.

[109] Leslie Ann Goldberg, Steven Kelk, and Mike Paterson. The complexity of choosing an H-colouring (nearly) uniformly at random. Technical Report ALCOMFT-TR-02-18, ALCOM-FT, Warwick, May 2002. Appearance: In *Symposium on the Theory of Computing*, 2002. to appear.

[110] Paul W. Goldberg. Estimating a Boolean Perceptron from its Average Satisfying Assignment: A Bound on the Precision Required. Technical Report ALCOMFT-TR-01-193, ALCOM-FT, Warwick, December 2001. Appearance: In *Proceedings of COLT 2001*, volume 2111 of *Lecture Notes in Artificial Intelligence*, pages 116–127. Springer-Verlag, 2001.

[111] Paul W. Goldberg. When Can Two Unsupervised Learners Achieve PAC Separation? Technical Report ALCOMFT-TR-01-192, ALCOM-FT, Warwick, December 2001. Appearance: In *Proceedings of COLT 2001*, volume 2111 of *Lecture Notes in Artificial Intelligence*, pages 303–319. Springer-Verlag, 2001.

[112] Marco Gori and Klaus Meer. A step towards a complexity theory for dynamical systems. Technical Report ALCOMFT-TR-01-191, ALCOM-FT, Aarhus, December 2001. Appearance: *Mathematical Logic Quarterly*, 2002. To appear.

[113] Fabrice Guillemin, Philippe Robert, and Bert Zwart. AIMD algorithms and exponential functionals. Technical Report ALCOMFT-TR-02-10, ALCOM-FT, INRIA, April 2002.

[114] Fabrice Guillemin, Philippe Robert, and Bert Zwart. Performance of TCP in the presence of correlated packet loss. Technical Report ALCOMFT-TR-02-85, ALCOM-FT, INRIA, May 2002. Appearance: In *15th ITC Specialist Seminar on Internet Traffic Engineering and Traffic Management*, Wurzburg, July 2002.

[115] Carsten Gutwenger, Michael Jünger, Sebastian Leipert, Petra Mutzel, Merijam Percan, and René Weiskircher. Advances in C-Planarity Testing of Clustered Graphs. Technical Report ALCOMFT-TR-02-127, ALCOM-FT, Cologne, May 2002.

[116] Carsten Gutwenger, Michael Jünger, Sebastian Leipert, Petra Mutzel, Merijam Percan, and René Weiskircher. Subgraph Induced Planar Connectivity Augmentation. Technical Report ALCOMFT-TR-02-126, ALCOM-FT, Cologne, May 2002.

[117] Alon Itai and Irit Katriel. Implicit dictionaries for internal and external memory. Technical Report ALCOMFT-TR-02-40, ALCOM-FT, MPI, May 2002.

[118] Predag Jelenkovic, Petar Momcilovic, and Bert Zwart. Reduced Load Equivalence under Subexponentiality. Technical Report ALCOMFT-TR-02-8, ALCOM-FT, INRIA, April 2002.

[119] Alexis Kaporis, Lefteris Kirousis, and Yannis Stamatiou. How to prove conditional randomness using the Principle of Deferred Decisions. Technical Report ALCOMFT-TR-02-75, ALCOM-FT, Patras, May 2002.

[120] Juha Kärkkäinen. Computing the Threshold for q-Gram Filters. Technical Report ALCOMFT-TR-02-43, ALCOM-FT, MPI, May 2002. Appearance: In *Proc. 8th Scandinavian Workshop on Algorithm Theory*, July 2002. To appear.

[121] Jan Klein, Jens Krokowski, Matthias Fischer, Rolf Wanka, and Friedhelm Meyer auf der Heide. The Randomized Sample Tree: A Data Structure for Externally Stored 3D-Scenes. Technical Report ALCOMFT-TR-02-64, ALCOM-FT, Paderborn, May 2002.

[122] Spyros Kontogiannis. Multiple Choice Assignments for Uniformly Related Machines. Technical Report ALCOMFT-TR-02-39, ALCOM-FT, MPI, May 2002. Appearance: In *34th ACM Symposium on Theory of Computing (STOC'02)*, page TO APPEAR, 2002.

[123] Arie M. C. A. Koster, Hans L. Bodlaender, and C. P. M. van Hoesel. Treewidth: Computational Experiments. Technical Report ALCOMFT-TR-02-63, ALCOM-FT, Utrecht, May 2002. Appearance: In Hajo Broersma, Ulrich Faigle, Johann Hurink, and Stefan Pickl, editors, *Electronic Notes in Discrete Mathematics*, volume 8. Elsevier Science Publishers, 2001.

[124] Dimitrios Koukopoulos, Marios Mavronicolas, Sotiris Nikoletseas, and Paul Spirakis. On the Stability of Compositions of Universally Stable, Greedy Contention-Resolution Protocols. Technical Report ALCOMFT-TR-02-74, ALCOM-FT, Patras and Cyprus, May 2002.

[125] Dimitrios Koukopoulos, Marios Mavronicolas, and Paul Spirakis. The Impact of Structure on Stability in Adversarial Queueing Theory. Technical Report ALCOMFT-TR-02-86, ALCOM-FT, Patras and Cyprus, May 2002.

[126] Dimitrios Koukopoulos, Sotiris Nikoletseas, and Paul Spirakis. Stability Issues in Heterogeneous and FIFO Networks under the Adversarial Queueing Model. Technical Report ALCOMFT-TR-02-100, ALCOM-FT, Patras, May 2002. Appearance: In *Proceedings of High Performance Computing 2001 – HiPC2001*, pages 3–14. Springer-Verlag, 2001.

[127] Elias Koutsoupias, Marios Mavronicolas, and Paul Spirakis. Approximate Equilibria and Ball Fusion. Technical Report ALCOMFT-TR-02-110, ALCOM-FT, Patras and Cyprus, May 2002. Appearance: In *9th International Colloquium on Structural Information and Communication Complexity (SIROCCO 2002)*. Carleton Scientific, 2002.

[128] Kim S. Larsen. Relaxed Red-Black Trees with Group Updates. Technical Report ALCOMFT-TR-01-186, ALCOM-FT, Aarhus, November 2001. Appearance: *Acta Informatica*. To appear.

[129] Insup Lee, Jin-Young Choi, Hee Hwan Kwak, Anna Philippou, and Oleg Sokolsky. A Family of Resource-Bound Real-time Process Algebras. Technical Report ALCOMFT-TR-02-139, ALCOM-FT, Cyprus, June 2002. Appearance: In M. Kim, B. Chin, S. Kang, and D. Lee, editors, *Proceedings of the 21st IFIP WG 6.1 International Conference on Formal Techniques for Networked and Distributed Systems (FORTE '01)*, pages 443–458. Kluwer Academic Publishers, 2001.

[130] Insup Lee, Anna Philippou, and Oleg Sokolsky. Formal Modeling and Analysis of Power-Aware Real-Time Systems. Technical Report ALCOMFT-TR-02-137, ALCOM-FT, Cyprus, June 2002. Appearance: *Computing and Control Engineering Journal*, To appear.

[131] Ulf Lorenz. Parallel Controlled Conspiracy Number Search. Technical Report ALCOMFT-TR-02-107, ALCOM-FT, Paderborn, May 2002. Appearance: In *Proceedings of the 13th Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA)*, pages 320–321, 2001.

[132] Ulf Lorenz and Burkhard Monien. The Secret of Selective Game Tree Search, when Using Random-Error Evaluations. Technical Report ALCOMFT-TR-02-108, ALCOM-FT, Paderborn, May 2002. Appearance: In H. Alt and A. Ferreira, editors, *Proceedings of 19th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 2285 of *LNCS*, pages 203–214. Springer, 2002.

[133] Thomas Lücking, Burkhard Monien, and Manuel Rode. On the Problem of Scheduling Flows on Distributed Networks. Technical Report ALCOMFT-TR-02-90, ALCOM-FT, Paderborn, May 2002. Appearance: In *27th International Symposium on Mathematical Foundations of Computer Science (MFCS)*. Springer LNCS, 2002. to appear.

[134] Bolette Ammitzbøll Madsen, Jesper Makholm Nielsen, and Bjarke Skjernaa. On the Number of Maximal Bipartite Subgraphs of a Graph. Technical Report ALCOMFT-TR-02-81, ALCOM-FT, Aarhus, May 2002.

[135] Conrado Martínez, Daniel Panario, and Alfredo Viola. Analysis of Quickfind with small subfiles. Technical Report ALCOMFT-TR-01-187, ALCOM-FT, Barcelona, November 2001. Appearance: In *Proc. 2nd Col. on Mathematics and Computer Science (MathInfo 2002)*, 2002. to appear.

[136] Kurt Mehlhorn and Ulrich Meyer. External-Memory Breadth-First Search with Sublinear I/O. Technical Report ALCOMFT-TR-02-48, ALCOM-FT, MPI, May 2002. Appearance: In *Proc. 10th Ann. European Symposium on Algorithms (ESA)*, LNCS. Springer, 2002, to appear.

[137] Mehlhorn, Kurt and Priebe, Volker and Schäfer, Guido and Sivadasan, Naveen. All-Pairs Shortest-Paths Computation in the Presence of Negative Cycles. Technical Report ALCOMFT-TR-02-42, ALCOM-FT, MPI, May 2002. Appearance: *Information Processing Letters*, 81(6):341–343, 2002.

[138] Ulrich Meyer. Buckets Strike Back: Improved Parallel Shortest-Paths. Technical Report ALCOMFT-TR-02-50, ALCOM-FT, MPI, May 2002. Appearance: In *Proc. 16th Intern. Parallel and Distributed Processing Symposium (IPDPS 2002)*. IEEE, 2002.

[139] Bengt Mueck, Wilhelm Dangelmaier, Matthias Fischer, and Wolfram Klemisch. Bidirectional Coupling of Simulation Tools with a Walkthrough System. Technical Report ALCOMFT-TR-02-65, ALCOM-FT, Paderborn, May 2002. Appearance: In *Proceedings of 'Simulation und Visualisierung'*, pages 71–84. SCS - The Society for Computer Simulation International, 2002.

[140] Jesper Makholm Nielsen. On the Number of Maximal Independent Sets in a Graph. Technical Report ALCOMFT-TR-02-87, ALCOM-FT, Aarhus, May 2002.

[141] Naomi Nishimura, Prabhakar Ragde, and Dimitrios M. Thilikos. Fast fixed-parameter tractable algorithms for nontrivial generalizations of vertex cover. Technical Report ALCOMFT-TR-02-111, ALCOM-FT, Barcelona, May 2002. Appearance: In *Proceedings Seventh International Workshop on Algorithms and Data Structures, WADS 2001, Springer Verlag, Lecture Notes in Computer Science, vol. 2125*, pages 75–86, 2001.

[142] Anna Philippou, Oleg Sokolsky, Insup Lee, Rance Cleaveland, and Scott Smolka. Hiding Resources that Can Fail: An Axiomatic Perspective. Technical Report ALCOMFT-TR-02-138, ALCOM-FT, Cyprus, June 2002. Appearance: *Journal of Information Processing Letters*, 80(1):1–13, 2001.

[143] Tomas Plachetka. Perfect Load-Balancing (Demand-Driven) Parallel Ray Tracing. Technical Report ALCOMFT-TR-02-95, ALCOM-FT, Paderborn, May 2002. Appearance: In B. Monien, editor, *Proceedings of the Euro-Par 2002*, LNCS. Springer, 2002. to appear.

[144] Tomas Plachetka, Olaf Schmidt, and Frank Albracht. The HiQoS Rendering System. Technical Report ALCOMFT-TR-02-94, ALCOM-FT, Paderborn, May 2002. APPEARANCE: In L. Pacholski and P. Ruzicka, editors, *Proceedings of the 28th Conference on Current Trends in Theory and Practice of Informatics (SOFSEM)*, volume 2234 of *LNCS*, pages 304–315. Springer, 2001.

[145] Tobias Polzin and Siavash Vahdati Daneshmand. Extending Reduction Techniques for the Steiner Tree Problem. Technical Report ALCOMFT-TR-02-135, ALCOM-FT, MPI, June 2002.

[146] Tobias Polzin and Siavash Vahdati Daneshmand. On Steiner Trees and Minimum Spanning Trees in Hypergraphs. Technical Report ALCOMFT-TR-02-34, ALCOM-FT, MPI, May 2002.

[147] Tobias Polzin and Siavash Vahdati Daneshmand. Partitioning Techniques for the Steiner Problem. Technical Report ALCOMFT-TR-02-38, ALCOM-FT, MPI, May 2002.

[148] Peter Sanders and Kurt Mehlhorn. Scanning Multiple Sequences Via Cache Memory. Technical Report ALCOMFT-TR-02-30, ALCOM-FT, MPI, May 2002. APPEARANCE: *Algorithmica*, 2002. to appear.

[149] Peter Sanders and Jesper Larsson Träff. The Hierarchical Factor Algorithm for All-to-all Communication. Technical Report ALCOMFT-TR-02-33, ALCOM-FT, MPI, May 2002. APPEARANCE: In *8th Europar*, LNCS. Springer, 2002. accepted.

[150] Peter Sanders and Berthold Vöcking. Random Arc Allocation and Applications. Technical Report ALCOMFT-TR-02-31, ALCOM-FT, MPI, May 2002. APPEARANCE: In *8th Scandinavian Workshop on Algorithm Theory*. Springer, 2002. to appear.

[151] Frank Schulz, Dorothea Wagner, and Christos Zaroliagis. Using Multi-Level Graphs for Timetable Information in Railway Systems. Technical Report ALCOMFT-TR-02-70, ALCOM-FT, Patras, May 2002. APPEARANCE: In *Algorithm Engineering and Experiments – ALENEX 2002*. Lect. Notes in Comp. Sc., 2002. To appear.

[152] Meinolf Sellmann and Warwick Harvey. Heuristic Constraint Propagation (Using Local Search for Incomplete Pruning and Domain Filtering of Redundant Constraints for the Social Golfer Problem). Technical Report ALCOMFT-TR-02-109, ALCOM-FT, Paderborn, May 2002. APPEARANCE: In *Proceedings CP-AI-OR*, pages 191–204, 2002.

[153] Meinolf Sellmann, Georg Kliewer, and Achim Koberstein. Capacitated Network Design - Cardinality Cuts and Coupled Variable Fixing Algorithms based on Lagrangian Relaxations. Technical Report ALCOMFT-TR-02-134, ALCOM-FT, Paderborn, June 2002. APPEARANCE: In *Proceedings of ESA'02, to appear*, 2002.

[154] Jop Sibeyn, James Abello, and Ulrich Meyer. Heuristics for Semi-External Depth First Search on Directed Graphs. Technical Report ALCOMFT-TR-02-49, ALCOM-FT, MPI, May 2002. APPEARANCE: In *Proc. Symposium on Parallel Algorithms and Architectures (SPAA)*. ACM, 2002. To appear.

[155] Anna Östlin and Rasmus Pagh. One-Probe Search. Technical Report ALCOMFT-TR-02-66, ALCOM-FT, Aarhus, May 2002. APPEARANCE: To appear in proceedings of ICALP 2002, 2002.

[156] Fariza Tahi, Manolo Gouy, and Mireille Régnier. Automatic RNA secondary structure prediction with a comparative approach. Technical Report ALCOMFT-TR-02-61, ALCOM-FT, INRIA, June 2002.

[157] Rolf Wanka. Any Load-Balancing Regimen for Evolving TreeComputations on Circulant-Graphs is Asymptotically Optimal. Technical Report ALCOMFT-TR-02-45, ALCOM-FT, Paderborn, May 2002. APPEARANCE: In *Proc. 28th Workshop on Graph-Theoretic Concepts in Computer Science (WG)*, 2002. to appear.

[158] Christos Zaroliagis. Implementations and Experimental Studies of Dynamic Graph Algorithms. Technical Report ALCOMFT-TR-02-69, ALCOM-FT, Patras, May 2002. APPEARANCE: In *Experimental Algorithmics - State of the Art*. Springer-Verlag, 2002. To appear.