Third Progress Report

ALCOM-FT

Algorithms and Complexity
Future Technologies

Project No. IST-1999-14186

January 2004

Summary

This is the third 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, 2002 to November 30, 2003.

ALCOM-FT brings together eleven of the leading groups in algorithms research in Europe in a project that aims at discovering new algorithmic concepts, identifying key algorithmic problems in important applications, and contributing 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.

The third period of the project was extended from twelve to eighteen months. This step was taken by the Consortium Board in agreement with the Commission to ensure proper finalization of some deliverables which had been delayed, mainly due to problems with filling programming positions. The step proved successful, as all deliverables have now been completed.

Besides this extension, the work has been carried out as planned. Proper coordination of the project was ensured by the efforts of the Consortium Board, the Work Package Leaders, and the Coordinator. Dissemination of the work done was ensured by the publication of scientific reports, by summer schools, by industrial talks, by specific dissemination oriented deliverables, and by the maintenance of websites for the entire project and for several deliverables within the project.

The deliverables completed during the third project period are listed in Table 2. The largest deliverable is D1, which represents a total of 206 scientific reports produced during the third project period. At least 146 of the 206 reports have already been published or been accepted for publication in the scientific community via key conferences and journals.

Deliverables D14, D15, D19, and D21 are the deliverables affected by the employment problems. In the contract amendment associated with the project extension, the delivery date of D14, D15, and D19 were moved from the second project year to the months stated (D21 was not moved, as it was only slightly delayed and had appeared by the time of the negotiation of the extension). These deliverables were completed on schedule according to their new dates. Deliverable D10, originally scheduled to appear in the first project year, is a volume in the *Lecture Notes in Computer*

Science series of *Springer Verlag* which had been delayed due to the time involved in the editorial process. Deliverable D10 was delivered in June 2003. The remaining deliverables were completed on schedule by the end of the project.

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 period 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.

No.	*Deliverable*	*Month*
D1	Research reports	Cont.
D10	Guidelines for algorithmic experiments (external release)	12
D14	PR web-pages for algorithmics	36
D15	External memory experimental platform (prototype)	36
D19	Production and transportation planning software prototype	36
D21	Testbed for experimental algorithmics (external release)	24
D23	Library of separation routines for ABACUS	36
D24	External memory experimental platform (final release)	42
D25	Distributed algorithmic engineering software package (final release)	42
D26	BSP-style library for dynamic distributed environments (final release)	42
D27	Production and transportation planning software user evaluation report	42
D28	Production and transportation planning problem instance database	42
D29	Complete software system for structured combinatorial optimization problems (SCIL)	42
D30	Testbed for experimental algorithmics (final release)	42
D31	Technology implementation plan	42
D32	Third progress report	42
D33	Final report	42

Table 2: The deliverables completed during the third project period.

Contents

Part I Work Packages

WP 1: Massive Data Sets

Participants: Aarhus, Barcelona, INRIA, MPI, Paderborn, CTI, Rome, Warwick
Coordinating site: Barcelona
Work Package Leader: José L. Balcázar
Number of technical reports: 30, of which 24 also appear in other work packages

Algorithms that process large datasets have to take into account that the cost of memory accesses depends on where the accessed data is stored. Traditional algorithm design is based on the von Neumann model which assumes uniform memory access costs. Actual machines increasingly deviate from this model. While waiting for a memory access, modern microprocessors can execute 1000 additions of registers. For hard disk accesses this factor can reach seven orders of magnitude.

ALCOM-FT has supported a *GI-Dagstuhl Research Seminar* on these topics organized as Dagstuhl-Seminar 02112 from March 10, 2002 to March 14, 2002 in the International Conference and Research Center for Computer Science at Schloss Dagstuhl. The seminar was held as a "self-taught" summer school where graduate students in cooperation with a few more experienced researchers have an opportunity to acquire knowledge about algorithms for memory hierarchies.

The main outgrowth of the seminar is a book published within Springer's LNCS series [179]. The 16 coherent chapters in this monograph-like tutorial book introduce and survey algorithmic techniques used to achieve high performance on memory hierarchies; emphasis is placed on methods interesting from a theoretical as well as important from a practical point of view. Peter Sanders has been invited to give talks on this subject at a German-Chinese workshop for theoretical computer science[1] and the SOFSEM conference 2004[2].

Whereas this major contribution is certain to have a profound impact on the research in the field, the project has contributed novel and important algorithms in several areas of this work package as described next. We concentrate on a few extremely important contributions, and wrap up with a brief survey of other contributions that may not seen as relevant as of now, although some of them have the potential to become a cornerstone of future developments, depending on the success of future research applications.

Algorithms tailored to extremely large datasets confront a number of difficulties, which, for the purpose of this presentation, we classify in two areas: those algorithms that effectively do allow for the use of the external memory devices in an efficient fashion, and those algorithms oriented towards some concrete, specific task while at the same time being aware of the fact that most of the data does not fit in internal memory. In turn, this last contribution opens the door to the analysis, not feasible up to now, of very large combinatorial structures, which we have started to perform as well. Along this line, we are also studying a number of problems that are popular in the Data Mining community, and intend to contribute to the progress of this discipline by helping to build more theoretical or soundly fundamental studies on which a later algorithm redesign task would be based.

Use of External Memory

The last level of the storage hierarchy in modern computers consists of reliable, persistent storage devices. Naturally, this level has the slowest access time, the largest capacity, and the cheapest costs.

[1] `http://www.iipl.fudan.edu.cn/activity/CG-ws.html`
[2] `http://www.cs.cas.cz/sofsem/04/design/`

Storage Networks

Nowadays, mainly magnetic hard disks are used to store large amounts of data permanently. But even the constantly growing capacity of disk drives cannot satisfy the ever growing demand for flexible and efficient storage systems. The simplest way to provide almost unlimited storage capacity is the use of disk drives in parallel. The underlying architecture of such a *storage network* may vary significantly. All it requires is a number of storage devices connected by an arbitrary network. Naturally, this network also provides communication links to the outside world.

Unfortunately, the efficient use of storage networks is not straightforward. There are a number of requirements, like data distribution, balance of data accesses, heterogeneity, or adaptivity, which have to be met to exploit the full potential of storage networks. In [193] we give an overview of existing techniques and algorithms addressing these problems. We start with classical solutions, like striping techniques and the use of redundancy, but the main focus of this work will be on more challenging tasks – namely heterogeneity and adaptivity. There are a number of solutions for each of these problems, but the most daring situations arise when all requirements must be met. We introduce some techniques that can handle many of these demands efficiently.

Some of the given techniques give rise to a more general view of storage networks and are capable of implementing a storage virtualization, i.e. the decoupling of the physical from the logical presentation of storage resources. In [48], we suggest new storage management concepts and we introduce a new management environment that is able to significantly reduce management costs and increases the performance and resource utilization of the given SAN infrastructure.

External Memory Platform

We have used our experience with algorithmic libraries, as explained in the initial proposal, to implement a first version of a platform for the experimentation with external memory algorithms, with a view to practical uses. The design of our external memory experimental platform was guided by two main goals: Very high performance and easy migration from existing software. The first objective implies support for parallel disks and overlapping of I/O and computation. Both objectives are well served by using C++ and templates. Finally, we observed that the C++ standard template library STL already supports most of the basic functionality we want to offer (iterators, streams, sorting, priority queues, search trees,...) so that it is possible to implement an external memory algorithms library with an interface that is part of the C++ language standard.

We started by configuring a state of the art Linux system that sustains full bandwidth I/O from up to eight disks and costs only 3000 Euro. Then we implemented a small I/O layer that supports both parallel disks and overlapping of I/O and computation using only Posix threads and ordinary synchronous file system I/O. The remaining library is independent of the operating system and hence highly portable. Our library `<STXXL>`[3] now has support for most of the algorithmic content of the STL except search trees which are under development in a diploma thesis project. The first academic outgrowth of the project was a new parallel disk sorting algorithm [93] that for the first time guarantees optimal overlapping of I/O and computation. Our implementation is also the first implementation of several recent innovations that for the first time guarantee close to optimal use of all disks. Our implementation is two to three times faster than previous algorithm libraries.

The `<STXXL>` library is now being used for several projects. Besides some of the graph algorithms described below, there is an ongoing project on suffix tree construction implementing algorithms described in [56, 152, 179] and three possible cooperations outside our group involving computer graphics, image processing, and geographic information systems.

Design and Analysis of Combinatorial Algorithms on External Memory

In [180] we present an I/O-efficient algorithm for Single-Source Shortest-Paths (SSSP) on undirected graphs $G = (V, E)$. Our algorithm needs $O(\sqrt{(VE/B)\log_2(W/w)} + \mathrm{sort}(V+E)\log\log(VB/E))$

[3] `http://www.mpi-sb.mpg.de/~rdementi/stxxl.html`

I/Os[4], where $w \in R^+$ and $W \in R^+$ are the minimal and maximal edge weights in G, respectively. For uniform random edge weights in $(0, 1]$, the expected I/O-complexity of our algorithm is $O(\sqrt{VE/B} + ((V+E)/B)\log_2 B + \text{sort}(V+E))$. The best previously known external-memory algorithm for SSSP required $\Omega(V)$ I/Os on general graphs. Hence, on sparse graphs with reasonably bounded weights our approach saves factors of nearly $\sqrt{B}$.

We are now using the `<STXXL>` library described above for implementing several of our theoretical algorithms. We already have an implementation for minimum spanning trees and connected components of massive graphs and there is an ongoing implementation project on external BFS.

Cache-oblivious Algorithms

In 1999, Frigo et al. introduced the concept of *cache oblivious* algorithms, which in an elegant way adapts the standard two-level memory model to the multi-level memory hierarchies found on actual computers. Cache-oblivious algorithms are algorithms optimized in the two-level model, except that one optimizes to an *unknown* block size B and memory size M. This simple change has significant consequences: Since the analysis holds for any block and memory size, it holds for *all* levels of the memory hierarchy simultaneously. 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, which increases portability.

In this period of the ALCOM-FT project, we have continued our investigations of the cache-oblivious model. A fundamental question raised by Frigo et al. is whether there is a *separation* in power between cache-oblivious algorithms and algorithms in the standard two-level model, i.e. a problem for which the asymptotical I/O complexity for all cache-oblivious algorithms is worse than for the best cache-aware. In [50], we resolve this open problem by showing a separation for the problems of *sorting* and *permuting*.

Specifically, we prove that I/O optimal cache-oblivious comparison based sorting is not possible without an assumption on the relation between M and B, and that there does not exist an I/O optimal cache-oblivious algorithm for permuting, not even in the presence of such an assumption. The lower bound proved for sorting is tight, since it matches the performance of the cache-oblivious sorting algorithm Lazy Funnelsort introduced in previous ALCOM-FT work.

In [28], we prove a separation for comparison based *searching*. In the two-level model, B-trees perform searches using $\log_B n$ I/Os, and cache-oblivious search trees exist which have asymptotically matching $O(\log_B n)$ complexity. However, we show that no such tree can do better that $c \cdot \log_B n$ for $c = \log_2(e) \approx 1.443$, and give a tree structure matching this for growing n and B.

In a more positive direction, we have investigated the practical feasibility of cache-oblivious sorting, and in [53] engineer a highly tuned version of Lazy Funnelsort, which we demonstrate competes very well with Quicksort. The more efficient use of memory makes it better than the fastest Quicksort we can find, already for instance sizes inside RAM (except on architectures with very fast RAM). This demonstrates that the overhead involved in being cache-oblivious can be small enough for the nice theoretical properties to transfer into practical advantages. Further details appears in the section on Work Package 5.

Trade-offs for External Memory Dictionaries

In terms of comparison complexity, the relationship between queries and updates in dictionaries is well-known, by the standard information theoretic lower bound and a result of Borodin et al. from 1981: Using only k comparisons on updates implies a lower bound on queries of $\max\{\log_2 n\,,\, n/2^k\}$, and this is asymptotically tight.

In terms of I/O complexity, B-trees perform queries and updates at $O(\log_B n)$ cost, which is best possible for queries, but not much is known about the trade-offs possible between query and update cost. In [49], we investigate this question, and provide lower bounds on the I/O cost for queries as functions of the I/O cost for insertions. We also describe data structures with query

[4] $\text{sort}(N) = \Theta((N/B)\log_{M/B}(N/B))$ is the I/O-complexity of sorting N data items.

and update costs asymptotically matching these lower bounds for large sections of the possible range of insertion cost.

Graph, Tree, and Hypergraph Algorithms

We continued to consider graph problems for large inputs.

The Webgraph is the graph whose nodes are the (static) HTML pages and (directed) edges are the hyperlinks between them. Today the size of the Webgraph is more than 3 billion pages. In [165] and [101] we present external and semi-external algorithms for computing properties of massive graphs, and for the large scale generation of graphs according to random webgraph models. We also report our findings on topological properties of such graphs together with the experiments on the time performance of the algorithms, and new properties of some stochastic graph models for the Webgraph presented in the literature. We have indeed applied our algorithms to the Webgraph by analyzing a crawl of about 200M pages and about 1.4 billion edges collected in 2001 by the WebBase project at Stanford. Specifically, by developing semi-external memory algorithms for computing disjoint bipartite cliques of small size, computing PageRank, computing strongly connected components, and for the large scale simulation of stochastic graph models, we were able to analyze the distribution of degree, the distribution of the size and of the number of strongly connected components, and more generally obtain an extensive study of the statistical properties of the Webgraph and similarly generated graphs. This last part is important to ensure that our results do not speak of a particular crawl in a particular year but reach a level of generality making them useful for the fast-evolving Webgraph.

Approximation and Sampling

The Alcom community proudly contributes to the development of algorithms with our knowledge of, and willingness to put to work, deep probabilistic algorithmic techniques leading to approximation algorithms or to guarantees of performance of sampling-based algorithms (see more in WP4). In this line, a major new result of the period is the LogLog Counting algorithm of [105] presented at ESA'03. The problem addressed is that of accurately estimating cardinalities, in particular the number of *distinct* records satisfying a given criterion. The algorithm proposed can estimate cardinalities in massive data sets of up to gigabytes of data in a single pass, using only one or two kilobytes of data, in which case the estimates it provides are normally within 1% or 2% of the actual value. Applications already exist for such algorithm, like detecting abnormal activity in routers, and further applications to data-mining are envisaged. The proposed algorithm is the best one currently known in terms of the accuracy produced for a fixed storage consumption. It can be fully distributed and lends itself to embedded hardware implementations. Some applications of the second year's work along the same line [136], of sampling to approximate values with statistical guarantees are described in the next subsection.

Applications of Sampling

Our analysis of sampling-based methods that adaptively decide whether the sample size is sufficient for a given approximation quality, instead of using a fixed sample size which is frequently inapplicable due to the too large size required by the standard Hoeffding-Chernoff analyses, was described thoroughly in [136]. Deliverable D16 was heavily based on that development, and has been extended beyond the initial specification; initially it was expected to consist of programs that use our adaptive sampling techniques to scale up three methods used in Machine Learning and in Data Mining, namely, finding frequent sets (like in the Apriori algorithm), constructing decision trees (like in the C4.5 software) and constructing weighted majorities of decision stumps by boosting. In the third year we have developed some additional applications, described below.

The standard Apriori algorithm does not allow direct implementation of these techniques, so we developed an alternative algorithm for the same task, described in [13] (and, in preliminary versions, in earlier reports). Additionally, we contributed a new analysis of the methods for

constructing association rules from the frequent sets [68]. Programs implementing these advances are included in deliverable D16. Recently, through a graduation project with participation of ALCOM researchers, we have developed an extension to a famous public-domain, open source, frequently used Data Mining package (Weka), adding to it the functionalities of our algorithm for boosted decision stumps based on adaptive sampling, plus an analogous improvement of the K-Means clustering algorithm. We believe that a good way of demonstrating the advantages of our approach is to integrate it into a tool that is widely used by practitioners. Thus, currently D16 contains four stand-alone programs (our algorithm for frequent sets as per [13], the scaled-up version of that algorithm as per [136], a program to experimentally evaluate association rules as per [68], and the scaled-up versions of the decision tree and decision stumps algorithms as per [136]) and two extensions of Weka (scaled-up weighted majorities of decision stumps and scaled-up K-Means). All the programs can be retrieved on the web from `http://www.lsi.upc.es/~gcasas/alcom/alcom_software.html`.

Glimpses of Additional Contributions

A number of advances made within the project have not been described in the above. For the sake of limiting the length of this report, we briefly describe here some of them, not even all: see the list of reports on the ALCOM-FT website for more information.

Multi-evaluation of the Coulomb potential induced by N particles is a central part of N-body simulations. The naive algorithm proceeds by summing over all pairs of particles exerting/experiencing force. This quadratic running time is however unacceptable for the very large systems of up to 10^5 objects to be simulated nowadays. In 3D, known sub-quadratic time algorithms return approximations up to given absolute precision. By combining data structures from Computational Geometry with fast polynomial arithmetic, we obtain in [210] approximations of prescribable relative error $e > 0$ within time $O(1/eN\text{polylog}N)$.

In [102] we introduce a random planted models of bi-categorical data to model the problem of collaborative filtering or categorical clustering, that is the problem in which, given a set of user preferences, the goal is to provide recommendations to users. We adapt the ideas of an algorithm due to Condon and Karp to develop a simple linear time algorithm to discover the underlying hidden structure of a graph generated according to the planted model with high probability. We also give applications to the probabilistic analysis of Latent Semantic Indexing (LSI) in the probabilistic corpus models introduced by Papadimitriou et al. Our experimental analysis shows that the algorithm might work quite well in practice. We show how to adapt the algorithm to work in secondary memory, and present the results of the experiments with massive data sets, i.e. graphs with more than 300M nodes.

With respect to the Data Mining applications, we have analyzed the statistical methods used to construct association rules from frequent sets [68], we have proposed a variation of the standard notion of episode mining [67], and we have concentrated on mining closed structures, which we see as having the potential to become the central notion in a more thoroughly developed theory of hypergraph-based data mining; namely, we have characterized exactly the association rules obtained from the closed-set approach in hypergraphs, in terms of tightest Horn approximations to propositional logic theories [12], and we are refining this approach to apply to the context of episode mining [69]. The advances are promising, and already within the few weeks after the end of the project we are obtaining consequences pertaining to propositional logic proper (as a semantic characterization of theories studied simply as syntactic calculi in the early eighties) and to computational learning theory (as new proofs of known algorithms that we are hoping to apply very soon to obtain algorithms for open problems in the area). In a closely related area, in [166] we present a fast algorithm for Information Retrieval in the vector model. The Latent Semantic Indexing (LSI) proved to be an effective technique in the field of Information Retrieval. Its drawbacks are the time needed to compute the SVD decomposition and to answer queries, since the query must be compared against each document in the collection, and this may not be feasible when we deal with massive data sets. In this paper we present a technique, called approximated LSI computation, that uses the information computed by the traditional LSI to provide a fast

online answer to the users. Informally, if compared to the traditional LSI, we can say that instead of returning the documents that are related to the concepts in the query, we return the documents that contain terms that are related to the concepts in your query.

A detailed study of the occurrences of generalized patterns in random texts is provided by [121]. The probabilistic model allows for unequal letter probabilities and for patterns having gaps. In this study, precise probability distributions are provided including explicit mean and variance estimates. This serves to discern signal from noise in large volumes of data as it makes it possible to determine which observed patterns are indicative of a behaviour deviating from expected statistical regularities. Potential applications are in the area of intrusion detection and the analysis of genomic sequences.

WP 2: Networks and Communication

Participants: INRIA, MPI, Paderborn, CTI, Rome, Warwick
Coordinating site: CTI
Work Package Leader: Christos Kaklamanis
Number of technical reports: 88, of which 58 also appear in other work packages

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 third year of the project addressed the following topics:

- Modelling of telecommunication networks,
- Wavelength optimization in optical networks,
- Algorithms for wireless networks,
- Virtual circuit and packet routing,
- Non-cooperative networks,
- Static and dynamic job allocation and load balancing, and
- Other investigations including network design, networked data structures, maintenance of dynamic networks, and hardware realization of networks.

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

Modelling of telecommunication networks

In this context, we aim at investigating the basic features of mathematical models describing telecommunication networks. The objective is firstly to give a precise description of the impact of the basic algorithms controlling the telecommunication networks and secondly to optimize the choice of their parameters. The equilibrium properties and transient behavior of the stochastic models of networks and telecommunication protocols are analyzed through Markovian methods. The renormalization (Euler's scaling) is the main fundamental tool to investigate the behavior of complex multi-dimensional Markov processes involved. The problems investigated include the following: (a) Stochastic modelling of TCP flows in ADSL traffic [27], (b) study of the interaction of long and short TCP flows in IP networks [46,143], and, (c) Connection Admission Control [104].

The aim of the study [27] is to characterize the Internet traffic in order to get realistic traffic models. For traffic analysis, a flow based approach and the popular mice and elephants dichotomy have been used: a mouse is by definition a TCP flow with less than 20 packets, the other TCP flows are elephants. Basically, due to its small size, a mouse is a TCP flow which do not use

the congestion avoidance algorithm of TCP to adapt its throughput to the traffic conditions. On the contrary, elephants share the remaining bandwidth with the flow control mechanism of TCP. These two types of flows have consequently a completely different behavior from a modelling point of view. Trace captures have been performed by France Telecom R&D on an Internet backbone link connecting different ADSL areas. The traffic is essentially commercial with a significant part of p2p applications.

After an extensive study of the statistics of the traces of the traffic, an approximation model has been proposed. Arrivals of mice have been represented as $M/G/\infty$, where customers are characterized by their arrival time, their service durations and their profile. Moreover, since the arrival rate is large, this process can be approximated by a Gaussian process. Elephants are represented as $M/G/1$ queues with the processor sharing discipline, because of the congestion on access links.

At the theoretical level, expressions for the Laplace transform of the stationary throughput and the stationary autocorrelation (including the transient behavior) have been obtained. In case of specific distributions of the transmission duration of the flows, exact asymptotics for the autocorrelation have been derived. These theoretical results have been partially validated with the ADSL traces.

We also investigate the performance of a network link supporting long elastic TCP flows, which may abort. The model consists of a network link crossed by elastic traffic with a transmission capacity assumed to be shared according to the Processor-Sharing discipline, which is the ideal bandwidth sharing achieved by TCP.

It is assumed that impatience of TCP flows occurs as follows: a flow aborts after some delay which is proportional to the total length of the transmitted file. In order to reduce the number of aborted flows, an admission control policy has been designed: A TCP flow may be rejected right just at the beginning of the connection. The policy is based on the number of TCP connections already using the link, it also depends on the statistics of the transfered files. The improvement of the global performance of the system with such a policy has been examined analytically, as well as the positive impact on the probability of aborting for a flow admitted in the network. Problems of fairness among flows have been also discussed: does a very long flow always abort even with this kind of admission control?

The basic queueing system considered in this analysis is the $M/G/1$ Processor-Sharing (PS) queue with or without impatience and with finite or infinite capacity. It is generally quite difficult to get explicit results concerning queues with impatience. For the $M/G/1$ PS queue under consideration, a result of independent interest has been proved: under some assumptions, a RSR (Reduced Service Rate) approximation (also called Reduced Load Equivalence) is shown to hold when the services are heavy tailed. This RSR result sheds some light on the fairness issues mentioned above and may be used to reduce the impatience of very large flows.

Wavelength optimization in optical networks

In the context of the wavelength assignment problem we consider two important network topologies: bi-directed trees and undirected rings.

For trees, one direction of research was to study greedy wavelength assignment algorithms that use randomization [66]. Such a randomized wavelength assignment algorithm can be thought of as applying some kind of randomized rounding on fractional path colorings. These are equivalent to solutions to the natural relaxation of the integer linear program corresponding to the wavelength assignment instance. In [66], we show that each symmetric set of paths of load L on a binary tree has a fractional path coloring of cost $1.367L$, where the load is defined as the maximum number of paths going through any fiber. The fractional colorings we construct have some nice properties which allow to obtain an $1.367 + o(1)$ approximation algorithm for wavelength assignment of symmetric sets of paths on binary trees (with some reasonable restrictions on the depth of the tree). Symmetric sets of paths are important since many services that will be supported by all-optical networks in the future will require bidirectional reservation of optical bandwidth.

Another direction of research was to try to approximate the optimal solution of the wavelength assignment problem using almost optimal solutions to the corresponding fractional path coloring problem as a guide. In [62], by simplifying and extending previous work, we present polynomial time algorithms that compute almost optimal fractional path colorings in bounded-degree trees and in rings. Our methods cover the case of multiple fibers as well. By applying a novel randomized rounding technique and using known wavelength assignment algorithms as subroutines, we obtain approximation algorithms with improved approximation ratios in bounded-degree trees and in rings with one or multiple parallel fibers. For the analysis, we use new tail inequalities for generalizations of occupancy problems.

Algorithms for wireless networks

Modelling and basic networks This line of research aims at the development and analysis of algorithms for network management and communication in mobile ad hoc networks (MANETs). In [181], we consider the problem of path selection in radio networks for a given static set of n sites in two- and three-dimensional space. For static point-to-point communication we defined measures for congestion, dilation, and energy consumption that take interferences among communication links into account.

We show that energy-optimal path selection can be computed in polynomial time and we introduce the diversity $g(V)$ of a set $V \subseteq \mathbb{R}^d$ for any constant d which can be used to upper bound the number of interfering edges. For real-world applications it can be regarded as $\Theta(\log n)$. A main result is that a c-spanner construction as a communication network allows to approximate the congestion-optimal path system by a factor of $O(g(V)^2)$. Furthermore, we show that there are vertex sets where only one of the performance parameters congestion, dilation, and energy can be optimized at a time. We show trade-offs lower bounding congestion $\times$ dilation and dilation $\times$ energy. The trade-off between congestion and dilation increases with switching from two-dimensional to three-dimensional space. For congestion and energy the situation is even worse. It is only possible to find a reasonable approximation for either congestion or energy minimization, while the other parameter is at least a polynomial factor worse than in the optimal network.

For the unidirectional transmission model where we allow to communicate in sectors in parallel we extended our investigations and measures. Graphs which can also be used to model such networks are the *random scaled sector graphs*. Their properties are studied in [98]. Other problems related to the modelling and basic operations in MANETs are studied in [41, 42, 77, 106, 186].

Dynamics and mobility In a next step, we allowed nodes to enter and leave the network and under these dynamics, we compared different network topologies for such basic networks, i.e. the *Yao-graph* and some also known related models, which are called the *SymmY-graph*, the *SparsY-graph*, and the *BoundY-graph* in [142]. We also present a promising network topology called the *HL-graph*. We compare them with respect to the degree and spanner-properties and we study their communication features. Our hardware model allows sector-independent directed communication, adjustable sending power, one frequency, and interference detection. We investigate how these network topologies bound the number of (uni- and bidirectional) interferences and whether these basic networks provide energy-optimal or congestion-minimal routing.

We also compare the ability of these topologies to handle dynamic changes of the network when radio stations appear and disappear. For this, we measure the number of involved radio stations and present distributed algorithms for repairing the network structure. It turned out that, in a worst-case scenario, the SparsY-graph combines good performance in terms of interferences, energy and congestion: For energy, it allows a constant factor approximation and a $O(\log n)$-approximation of the congestion. However, all Yao-graph based topologies have only linear time algorithm for rebuilding the graph after one station appears or disappears, because a linear number of stations is involved in the worst case. For the HL-graph we need only logarithmic time and a logarithmic number of involved stations. Furthermore, it provides a linear approximation of the energy optimal path system, and allows path systems approximating the minimal congested routing by a factor of $O(\log^2 n)$.

Additionally to this kind of dynamics, participants may move freely and unpredictably. In [198] we investigate distributed algorithms for mobile ad hoc networks for moving radio stations with adjustable transmission power in a worst case scenario. We consider two models to find a reasonable restriction on the worst-case mobility: the *pedestrian* and the *vehicular* model. Our goal has been to maintain persistent routes with nice communication network properties like hop-distance, energy-consumption, congestion and number of interferences. A route is persistent, if we can guarantee that all edges of this route can be uphold for a given time span Δ, which is a parameter denoting the minimum time the mobile network needs to adopt changes, i.e. update routing tables, change directory entries, etc. This Δ can be used as the length of an update interval for a proactive routing scheme.

We extend some known notions such as transmission range, interferences, spanner, power spanner and congestion to both mobility models and introduced a new parameter called crowdedness that states a lower bound on the number of radio interferences. Then we proved that a mobile spanner hosts a path system that polylogarithmically approximates the optimal congestion. We present distributed algorithms which construct mobile spanners with low congestion, low interference number, low energy-consumption, and low degree. We measure the optimality of the output of our algorithm by comparing it with the optimal choice of persistent routes under the same circumstances with respect to pedestrian or vehicular worst-case movements, and finally, we present solutions for dynamic position information management under our mobility models.

In [100], we study an important combinatorial notion, the *infection time* of graphs, by using concurrent random walks of k particles, one of which is red and infects the rest $k-1$ white particles when meeting them on the graph. We develop a set of probabilistic tools that are used to obtain tight upper bounds for general graphs and important special cases (cliques, expanders). This notion may be useful in estimating the information propagation time in ad-hoc mobile networks.

Energy-efficient communication In [130] we continue our work on energy efficient routing in radio networks using advanced data structures from algorithmic geometry. Suppose we have n radio stations whose position is known and we assume that the energy consumption for communication between two stations is bounded by their Euclidean distance raised to some constant. In this situation, the energy consumption can be reduced by using intermediate stations as relay stations. We develop a linear space data structure that can be precomputed in time $O(n \log n)$ and allows constant time queries for k-hop paths whose energy consumption is within a factor $1 + \epsilon$ from optimal for any constant $\epsilon > 0$.

In [64], we address the issue of supporting typical communication patterns like broadcasting, multicasting, and gossiping in wireless networks using the minimum amount of energy. We consider a series of *wireless network design problems* which are formulated as follows. Given a wireless network with omnidirectional transmitters modelled as a complete directed graph $G = (V, E)$, where $|V| = n$, with a non-negative edge cost function $c : E \to R^+$, and a non-negative node weight assignment $w : V \to R^+$, the *transmission graph* G_w is the directed graph having the same set of nodes as G and the directed edges (u, v) only if the weight assigned to node u is at least the cost of the edge (u, v), i.e., $w(u) \geq c(u, v)$. Intuitively, the weight assignment corresponds to the energy levels at which each node operates while the cost between two nodes indicates the minimum energy level necessary to send messages from one node to the other. Usually, the edge cost function is symmetric (i.e., $c(u, v) = c(v, u)$). Asymmetric edge cost functions can be used to model medium abnormalities or batteries with different energy levels. The problems we study can be generally stated as computing an energy assignment to the nodes of the networks so that the transmission graph maintains a connectivity property and the sum of weights is minimized.

We consider the problems of minimum-energy broadcasting, multicasting, and gossiping in networks with symmetric and asymmetric edge-cost functions. In asymmetric graphs, we obtain optimal logarithmic approximations for broadcasting and gossiping, and polylogarithmic inapproximability results for multicasting. For networks with symmetric edge-cost functions, we obtain logarithmic approximations for multicasting and constant approximations (and inapproximability results) for group communication problems. Our results exploit the relations of these wireless net-

work design problems to famous combinatorial problems like steiner forest, node-weighted steiner tree, node-weighted connected dominating set, directed steiner tree, set cover, and more.

Frequency assignment Frequency assignment is important in wireless networks utilizing frequency division multiplexing (FDM) technology. These may be ad hoc networks consisting of transmitters which can communicate directly among them, or networks of base stations aiming at supporting communication for mobile phone users. FDM assigns different frequencies to users or transmitters so that users or transmitters which are close to each other are assigned different frequencies. Since, usually, the spectrum of available frequencies is limited, it is important to minimize the number of frequencies used, or given the available spectrum, to maximize the number of connections.

These problems can be formulated as coloring (also known as radio labeling and radiocoloring) or independent set problems in graphs. In this context, in [184] we employ the Probabilistic Method to the radiocoloring problem (RCP). In particular, we study the optimization version of RCP where the objective is to minimize the number of colors used. We first provide an upper bound for the minimum number of colors needed to radiocolor a graph of girth at most 7. Then, we study whether the minimum order of a radiocoloring assignment is determined by local conditions, i.e. by the minimum order radiocoloring assignment of some small subgraphs of it, by discussing relevant theorems in the literature. We also study another version of the problem in [34] where some nodes may have a pre-assigned frequency, and obtain several algorithmic, combinatorial, and complexity-theoretic results. In another line of relevant research, our approach to problems which are intractable on general networks is to look for efficient algorithms for specialized inputs, i.e., when we may assume that the input graph is of a special type. Using this approach, several basic problems including coloring and independent set are studied and proved to be polynomial-time solvable on graphs that have no gem or path of length five as induced subgraph [47].

In [6] we provide a new coloring algorithm which colors the square of any planar graph G using at most $1.5\Delta(G) + c$ colors, where $c \leq 26$, and the maximum degree of G (i.e. $\Delta(G)$) is greater than 8. In 1977, Wegner conjectured the above bound for $c = 1$. The best previously known upper bound for the chromatic number of the square of G was $1.66\Delta(G) + 24$.

Continuing previous work, we consider an online version of the independent set problem in cellular networks in [65]. These are planar networks with specific structure modelling interferences between base stations with equivalent transmitters uniformly distributed on the plane. We present improved competitive ratios and competitiveness lower bounds for the problem of approximating the maximum independent set when the network nodes (representing users wishing to connect to base stations) appear on-line. Our algorithms are very simple and make competitive decisions using only a constant number of random bits or comparable weak random sources.

Virtual circuit and packet routing

Virtual circuit routing In the virtual circuit routing problem, we are given a graph and a sequence of routing requests between node-pairs in the graph. The goal is to connect all pairs via a path in the network such that the congestion (i.e., maximum traffic on a network link) is minimized. Previous work on this problem mainly focuses on the online or offline version of this problem, where a routing decision may depend on other requests.

In [189], we have shown that for undirected networks there exist also an efficient *oblivious* algorithm for this problem, i.e., an algorithm that does not use any knowledge of the request sequence for its routing decisions. Such an algorithm can be implemented very efficiently in a distributed environment since it only requires static routing tables at the network nodes. Our algorithm achieves a competitive ratio of $O(\log^3 n)$ with respect to the congestion of the network links. The algorithm is based on a hierarchical decomposition of the network with certain properties. It is shown in [189] that such a decomposition exists and can be constructed by a non-polynomial time algorithm. Therefore the result is non-constructive. In [31], we refine the decomposition algorithm of [189] to get a polynomial time version that creates a hierarchical decomposition that guarantees a competitive ratio of $O(\log^4 n)$. Independently, Harrelson et al. (SPAA 2003) also

obtained a constructive version of the decomposition algorithm that even leads to a competitive ratio of $O(\log^2 n \log\log n)$.

Bufferless routing In this area, we are interested in investigating the possibility or impossibility, and the corresponding costs, of devising *packet-routing algorithms* for communication networks that achieve nice performance guarantees with respect to the *congestion* C (the maximum number of packets that simultaneously cross a *link* of the network, and the *dilation* D (the maximum number of links a packet needs to traverse from *source* to *destination*). A well known lower bound of $\Omega(C+D)$ sets up a measurable challenge for the performance of any efficient routing algorithm: when is it possible to devise routing algorithms with *Routing Time* (the time for the last packet to reach its destination) attaining this lower bound? In [58], we have embarked on a study of this fundamental question in the context of *bufferless* routing algorithms (algorithms that may not use *buffers* for intermediate storage of packets). In particular, we have considered two restrictions of *Bufferless Routing Algorithms,* namely *Hot-Potato Routing Algorithms* and *Direct Routing Algorithms*.

In *hot-potato* routing, nodes in the network have no buffers for packets in transit, and packets may instead be deflected. A hot-potato routing algorithm is *greedy* if packets are advanced from their sources towards their destinations whenever possible. We consider greedy, hot-potato routing algorithms for the specific case of a *tree* with n nodes. We have devised two greedy, hot-potato routing algorithms: (a) A deterministic algorithm, which has a routing time of $O((\delta C + D)\lg n)$, where δ is the maximum node degree, and (b) a randomized algorithm, which has a routing time of $O((C+D)\lg^2 n)$ with high probability. These are the *first* known hot-potato routing algorithms for trees whose routing time is within logarithmic factors of the $\Omega(C + D)$ lower bound.

Direct Routing is a special case of bufferless routing in which N packets must be routed *without conflicts* along specific paths. In the report [57], we first investigate the computational complexity of *Direct Routing*. By reduction from *Vertex Coloring*, we prove that the Direct Routing problem is $\mathcal{NP}$-complete, and that it is as hard to approximate as Vertex Coloring. This implies that buffering is necessary in order to tractably compute direct routing schedules. Further on, we constructed a hard routing problem for which any direct routing algorithm requires a vast amount of total packet buffering to *approximate* an optimal routing schedule with good approximation ratio.

We have also studied algorithms for efficient $(\widetilde{O}(C+D))$ Direct Routing. We devised a generic, greedy direct routing algorithm for general direct routing problems with direct routing time that is worst-case optimal. Finally, we have considered as case studies many interesting routing problems on commonly used network topologies such as trees, meshes, hypercubes and butterflies.

Adversarial packet-switching networks In this area, we continue our study of *stability* and *instability* properties of *packet-switching networks* within the framework of *Adversarial Queueing Theory,* which assumes that packets are adversarially injected into the network. We remain interested in understanding which *protocols* remain *stable* on which *networks* and under which *model assumption*. Recall that a packet-switching network is *stable* if the number of packets in the network remains bounded at all times.

We have addressed the very natural question of how network structure precisely affects the stability properties of networks [157]. We consider *size, diameter, maximum vertex degree,* minimum number of *disjoint paths* that cover all edges of the network, and *network subgraphs* as crucial structural parameters of the network. We have assembled a comprehensive collection of structural results, in the form of stability and instability bounds on injection rate of the *adversary,* for various greedy protocols such as FIFO.

We have also investigated in [158, 159] the effect of dynamically changing the *capacities* of a network on their stability and instability properties. To give an essence of the questions that we have considered, think of whether dynamic capacities could turn a *universally stable protocol,* that is a protocol that remains stable on all networks if capacities are static, to an unstable one on some particular network as a result of dynamically changing the capacities. Also, for a network

that we know is unstable with a certain protocol running on it and with capacities constant, how is its *instability threshold* affected when capacities are now allowed to change dynamically? In this line of research, we assume that each link capacity may take on integer values from the interval $[1, C]$ with $C > 1$, and we assume a particular adversary called a (w, ρ)-*adversary*.

We have managed to prove *tight* lower bounds on the *instability thresholds* of several natural, greedy protocols, such as LIS (*Longest-In-System*), SIS (*Shortest-In-System*), and *compositions* of them. We have also proved that any arbitrary network G (possibly *cyclic*) running a greedy protocol remains stable for a rate not exceeding a particular *stability threshold*; the stability threshold is a function of C and the length of the longest path in the network. We also consider the impact of dynamically changing link capacities on the performance bounds of LIS (*Longest-In-System*) and SIS (*Shortest-In-System*) protocols on *Directed Acyclic Graphs* (DAGs). Although it is known that any greedy protocol is stable on a DAG, this does not guarantee any reasonable bounds on the packet delay. Our study has established that the packet delays for LIS on $DAGs$ are $\Theta(iw\rho C)$, where i is the length of the longest path leading to a node when nodes are ordered by the topological order induced by the DAG. We also prove a corresponding lower bound for SIS.

Non-cooperative networks

Our work in this area has considered three different models for a non-cooperative network. These are the *KP Model,* the *Hybrid Model* and the *Restricted Links Model*. In the KP model, a single source and a single destination are connected by m parallel links. Each of n users chooses a strategy, that is, a probability distribution over all links, trying to minimize its private cost to send its traffic from source to destination. The private cost is defined as the maximum expected latency over all links the user has chosen with positive probability. This definition of private cost often contradicts the social cost, defined as maximum expected latency on a link, which measures the global performance of the network. A Nash equilibrium is a state of the system such that no user can decrease its private cost by unilaterally changing its strategy. For a survey on results in the KP-model we refer to [116].

Computing the best-case pure Nash equilibrium with respect to the social cost is NP-hard. Providing algorithms to compute a pure Nash equilibrium from a given schedule without increasing the social cost, we get a polynomial-time approximation scheme (PTAS) to approximate a best-case pure Nash equilibrium for the model of identical links [133] and related links [115], respectively.

In [133], we give structural results providing substantial evidence for the Fully Mixed Nash Equilibrium Conjecture, henceforth abbreviated as FMNE Conjecture, which states that the worst Nash equilibrium with respect to the social cost is the fully mixed Nash equilibrium, where each user chooses each link with positive probability. Specifically, we prove that the FMNE Conjecture is valid for pure Nash equilibria and that under a certain condition, the social cost of any Nash equilibrium is within a factor of $6 + \varepsilon$, of that of the fully mixed Nash equilibrium, assuming that link capacities are identical. In [172], we prove the FMNE Conjecture for special instances and show that it does not hold for $n = 3$ and $m = 2$ in case of unrelated links.

In [171] we consider the Hybrid model (which combines in an interesting way features of the KP model and the *Wardrop* model), and we evaluate the performance of Nash equilibria with the help of *quadratic social cost* as the sum of the expectations of the squares of the incurred link latencies. We present a comprehensive collection of combinatorial expressions for the computation of quadratic social cost. These expressions readily imply some polynomial algorithms to compute quadratic social cost in several special cases, and a corresponding *general,* pseudopolynomial, dynamic programming algorithm. Furthermore, we prove that, for the model of identical users and identical links, the worst-case Nash equilibrium is the fully mixed Nash equilibrium. Finally, we present a comprehensive collection of (usually tight) constant (that is, independent of m and n) lower and upper bounds on the coordination ratio.

In [132] we consider the *Restricted Links Model* which is the variant of the KP model in which each user may only be routed on a link from a certain set of *allowed links* for the user. Our study of the Restricted Links Model has led us to the development of new techniques for the computation of Nash equilibria based on the recently studied *unsplittable flows* and successfully employ the

famous *Preflow-Push* algorithm in the setting of unsplittable flows. We have derived a polynomial algorithm to compute, from any given assignment, a pure Nash Equilibrium with non-increased Maximum Social Cost. We have relied very strongly on ideas from *blocking flows*. and we use techniques similar to those in the generic *Preflow-Push* algorithm to approximate a schedule with minimum Social Cost, gaining an improved approximation factor. We extend this result to the model of related links, gaining an approximation factor of 2. Both our approximation algorithms run in polynomial time. Finally, we have also managed to prove tight bounds on Coordination Ratio (restricted to pure Nash Equilibria) for the Restricted Links Model.

In [24], we study algorithmic questions concerning a congestion game in which there is a single provider that offers a service to a set of potential customers. Each customer has a particular demand of service and the behavior of the customers is determined by utility functions that are non-increasing in the congestion. Customers decide whether to join or leave the service based on the experienced congestion of the offered prices. If the prices of services are fixed, then the customer behavior leads to a pure, not necessarily unique Nash equilibrium. In order to evaluate marketing strategies, the service provider is interested in estimating its revenue under the best and worst customer equilibria. For the case in which the provider has perfect knowledge of the behavior of all customers, we present a complete characterization of the *complexity of computing optimal pricing strategies* and of *computing best and worst equilibria*. Basically, we show that most of these problems are inapproximable in the worst case but admit an "average-case FPAS". We extend our analysis to a more realistic model in which the provider has *incomplete knowledge* and show that the worst-case complexity of the problem increases with the precision of the available knowledge.

Static and dynamic job allocation and load-balancing

Graph-partitioning and load-balancing In order to develop and test efficient graph partitioning heuristics, we have computed new lower bounds on the cut-size of graph partitions. In [110], we show that the existing spectral bounds are tight for a few graphs only, and improve these bounds by a quadratic factor for grid-like structures. Applying similar methods, we compute new spectral bounds for the edge- and node-expansion of graphs in [112]. These new bounds perform well for graphs having large expansion and bounded vertex degree. In addition, we generalize these bounds to edge-weighted graphs.

In the same paper [112], we also consider load balancing schemes for networks containing indivisible unit size tokens. In some previous papers, it has been shown that classical diffusion algorithms do not achieve a good balanced state. We applied a randomized approach based on diffusion, and have shown that, using this method, we obtain an asymptotically optimal balanced situation. In [111], we analyze the performance of several diffusion based load balancing strategies in the case of arbitrary divisible load items. In [109], we construct scalable network topologies having a bounded vertex degree and a small number of different eigenvalues. We have shown that these topologies are well suited for load balancing applications and can be used in parallel or distributed networks to bound the communication overhead for such applications.

In the report [97] we give the best bounds for the bisection width of 3- and 4-regular graphs improving previously known bounds. This result is generalized in [83] where the case of d-regular graphs is considered.

Scheduling malleable tasks In [88] we consider the problem of finding a schedule for n identical malleable tasks on p identical processors with minimal completion time. This problem arises while using the branch & bound or the divide & conquer strategy to solve a problem on a parallel system. If nothing is known about the sub-problems, then they are assumed to be identical. We assume that the execution time decreases with the number of processors while the computational work increases. We give an algorithm with running time exponential in p which computes an optimal schedule. In order to approximate an optimal schedule, we use the concept of phase-by-phase schedules. Here schedules consist of phases in which every job uses the same number of processors. We prove that one can approximate an optimal schedule up to a factor 5/4 using

constant time. We show that this is optimal. Furthermore, we give an ε-approximation algorithm if the speed-up is optimal up to a constant factor.

Online scheduling In [21–23] we consider the problem of efficiently serving on-line requests submitted by interactive users to a distributed server system. This setting is modelled in a very general way, as the problem of on-line scheduling jobs on a set of identical machines to minimize some function of time.

In [23] we review work done on scheduling to minimize average response time or related metrics, on single and parallel machines, considering the classical abstract model, in which a set of jobs is presented on line to a set of identical machines. Each job has a processing time and has to be processed, possibly over a non-continuous interval, for an overall amount of time equal to its processing time. The response time of a jobs is the difference between its completion and release times, while the stretch is the ratio between the response time of a job and its size. Average response time and average stretch are defined accordingly.

In [21] we study the quality of service (QoS) that is achievable by semi-clairvoyant online scheduling algorithms, which are algorithms that only require approximate knowledge of the initial processing time of each job, on a single machine. It was an open question whether it was possible for a semi-clairvoyant algorithm to be O(1)-competitive with respect to average flow time on one single machine. The authors settle this open question by giving a semi-clairvoyant algorithm that is O(1)-competitive with respect to average flow time on one single machine. They also show a semi-clairvoyant algorithm on parallel machines that achieves up to constant factors the best known competitive ratio for clairvoyant on-line algorithms. It is known that the clairvoyant algorithm SRPT is optimal with respect to average flow time and is 2-competitive with respect to average stretch. Thus it is possible for a clairvoyant algorithm to be simultaneously competitive in both average flow time and average stretch. In contrast they show that no semi-clairvoyant algorithm can be simultaneously O(1)-competitive with respect to average stretch and O(1)-competitive with respect to average flow time.

In [22] we introduce the notion of smoothed competitive analysis of online algorithms. Smoothed analysis has been proposed by Spielman and Teng to explain the behaviour of algorithms that work well in practice while performing very poorly from a worst case analysis point of view. We apply this notion to analyze the Multi-Level Feedback (MLF) algorithm to minimize the total flow time on a sequence of jobs released over time when the processing time of a job is only known at time of completion. A direct consequence of our result is also the first average case analysis of MLF. We show a constant expected ratio of the total flow time of MLF to the optimum under several distributions including the uniform distribution.

Other investigations

Efficient network design and cost allocation In [167], we study the problem of efficiently connecting selfish users to facilities, devising a cost sharing mechanism, such that the optimal cost of connecting users to the facilities is recovered, and the users do not pay more than this cost. We present cost sharing methods for connected facility location games that are cross monotonic, competitive, and recover a constant fraction of the cost of the constructed solution. The novelty of this paper is that we use randomized algorithms, and that the expected cost is shared among the participating users. As a consequence, these cost sharing methods are simple, and achieve attractive approximation ratios. We also provide a primal-dual cost sharing method for the connected facility location game with opening costs.

Networked data structures During the third year of the project, we have embarked on a study [59] of the possibility or impossibility, and the corresponding costs, of devising *concurrent, low-contention* implementations of atomic Read&Modify&Write (or RMW) *operations* in a distributed system. We considered here a natural class of *monotone* RMW operations associated with *monotone groups,* a certain class of algebraic groups that we specifically introduced. The popu-

lar Fetch&Add and Fetch&Multiply operations are examples from the class of monotone RMW operations we considered.

We devised, as our chief combinatorial instrument, a *Monotone Linearizability Lemma,* which establishes inherent ordering constraints of *linearizability* for a certain class of executions of *any* distributed system that implements a monotone RMW operation. The end results of our study specifically apply to implementations of (monotone) RMW operations that are based on *switching networks*, a recent class of concurrent, low-contention data structures that generalize *counting networks* (which implemented the Fetch&Increment operation). These results are negative; they are shown through the *Monotone Linearizability Lemma*.

In particular, we have shown that if the network incurs low contention, then its size must be infinite. For switches with a finite number of states, *bounded* concurrency (number of concurrent processes) already suffices to prove the lower bound, while a separate proof employs *unbounded* concurrency in order to compensate for switches with an infinite number of states. Also, we show that any such network has *sequential* executions with n processes and latency $\Omega(n)$. Since it is known that Fetch&Increment is implementable with counting networks of *finite* size and of latency $\Theta(\lg n)$, our lower bounds imply a time and space complexity separation between Fetch&Increment and any monotone RMW operation in the model of switching networks. Our lower bounds provide a mathematical explanation for the observed inability of researchers over the last twelve years or so to extend counting networks, while keeping them finite-sized, high-concurrency, low-latency and low-contention, in order to perform tasks more complex than just incrementing a counter by one, but yet as simple as adding an arbitrary value to a counter.

Maintenance of dynamic networks In [11], we give a deterministic algorithm for maintaining a t-spanner of a general weighted undirected graph under a sequence of edge insertions and deletions. A t-spanner S of an undirected graph G is a subgraph of G such that for any pair of vertices a, b we have $d_S(a, b) \leq t \cdot d_G(a, b)$, where d_S and d_G are respectively the distances on S and G. In worst-case sequences our algorithm requires $\widetilde{O}(t \cdot n2 \cdot \log C)$ amortized time per update, where n is the number of vertices and edge costs are real numbers in $[1, C]$ (the $\widetilde{O}$ notation hides polylogarithmic factors). The spanner has $O(t \cdot n^{1+4/(t+2)} \cdot \log C)$ edges. For random update sequences, the average cost per update is improved to $\widetilde{O}(t \cdot n^{1+4/(t+2)} \cdot \log C)$.

Hardware realization of networks For visual display and for hardware realization of networks in several contexts, it may be required to find an embedding of the network on the plane, with edges only horizontal or vertical line segments. A new characterization of the graphs with this property was obtained and it was shown that a similar characterization does not exist when we require that each angle is a multiple of $2\pi/d$ for $d = 3$ or $d > 4$ [38].

WP 3: Production and Transportation Planning

Participants: Cologne, Paderborn, MPI, Utrecht
Coordinating site: Utrecht
Work Package Leader: J.A. Hoogeveen and Jan van Leeuwen
Number of technical reports: 12, of which 6 also appear in other work packages

WP 3 aims to explore novel 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 that arise in industry and which require both practical algorithm design and fundamental research.

Below we describe the work done in WP3 in the last one and a half 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 [129], we studied the problem of taking reservations for transportation systems with many stops where people can enter or exit. The problem situation can be described as follows. On a first-come-first-served basis a customer specifies his starting and ending station, after which the passenger is accepted, that is, he is told that he will be assigned a seat, or rejected. The decision of acceptance/rejection is based on whether it is possible to seat this passenger together with all previously accepted passengers. In this model, the actual seat assignment occurs off-line, after the whole set of accepted passengers is known; the seat is the same for the whole trip. We studied the acceptance/rejection problem and solved this problem optimally through an algorithm with optimal time complexity.

We continued our work on planning problems. One of these concerns the integration of different optimization problems for the airline industry. Together with industrial partners we integrate the Airline-Fleet-Assignment and Aircraft-Rotation-Building phases. Thereby the resulting fleet assignments represent a better starting point for the following planning phases of an airline.

In [141] the classical Fleet-Assignment problem, in which it has to be decided which plane (or type of plane) is used to perform a flight, is extended by the capability to handle rotational constraints. This type of constraints, one example of which are the so-called connection-dependent ground-times, can be used to model minimum delays between two flights depending on both the incoming and the outgoing flight. Three different solution methods are presented, one MIP- and two Local-Search-based approaches, which we evaluate empirically on a number of real-world problem instances.

Furthermore, we continued our research on the home health care problem in close co-operation with industrial partners. Home health care (HHC), i.e. visiting and nursing patients in their homes, is a growing sector in the medical service business. From a staff rostering point of view, the problem is to find a feasible working plan for all nurses that has to respect a variety of hard and soft constraints and preferences. Additionally, home health care problems contain a routing component: A nurse must be able to visit her patients in a given roster using a car or public transport. It is desired to design rosters that take both the staff rostering and vehicle routing components in consideration while minimizing transportation costs and maximizing satisfaction of patients and nurses. In [30] we present a compact model for the home health care problem which is flexible enough to break down most real-world HHC problems of different characteristics. Especially, in [30] we describe the core optimization components of the software that we have produced. In the optimization kernel, a combination of linear programming (LP), constraint programming (CP), and (meta-)heuristics for the home health care problem is used, and we show how to apply these different heuristics efficiently to solve home health care problems. The overall concept can be adapted to various changes in the constraint structure, thus providing the flexibility needed in a generic tool for real-world settings. The industrial software-prototype is currently undergoing intensive user evaluation, where solutions generated by our optimization methods are competing in real-world scenarios. In a further step, requests for additional legal and company constraints, or credit point systems (contributed by researchers from ergonomics) will be integrated into our algorithmic framework.

Another line of work concerns the Capacitated Network Design Problem (see also [200] from the second project period), 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 developed system consists of several algorithmic components, which can be combined into an exact or a heuristic branch-and-bound algorithm. The Lagrangian relaxation is solved by subgradient search or by using a bundle-based method. We have extended our system to a branch-and-cut system by proposing several problem-specific cuts. The cuts are generated dynamically and added to the Lagrangian objective function when they are violated. The performance of the implemented system components is compared and evaluated on a large set of benchmark instances. Our system gives high quality results and outperforms currently known systems in running time and solution quality. The results were partially published on the ESA 2002 and at the international symposium on Mathematical Programming (ISMP 2003).

In our work on scheduling problems, one of the problems concerns the no-wait job shop problem, which is used to model the production planning problem of a pharmaceutical company. Here the different steps in a production process are connected and between every two consecutive steps there is a fixed time delay. Hence, when the starting time of the first step of a process is known, then the time-path of the whole process is known. The goal is to minimize the time by which all production processes have been completed. We have developed a method to solve this kind of problems that is based on integer linear programming and constraint programming; this has been implemented in a prototype (deliverable D19). The prototype will be improved in the next couple of months by incorporating the lower bounds described in [206] (see below for a short description) and the lower bound of the LP-relaxation for the integer linear programming formulation using time-indexed variables. A report on the connection between the no-wait job shop problem and the asymmetric traveling salesman problem is in preparation.

We also continued our work on the problem of accepting or rejecting jobs in machine scheduling; if a job is accepted, then it must be completed at a given deadline to earn a (unit) reward. Processing job j results in an amount of p_j units of work to be done, where p_j can be either deterministic or stochastic, which must occur on one of the machines, which can either have a deterministic or stochastic output. In a previous report [2] we have discussed the variant with one machine with deterministic output and stochastic variables p_j. In [3], we discuss the case with deterministic p_j values and and machines with deterministic output and some side-constraints. Currently we are working on the case with machines with stochastic output, and p_j values that are either deterministic or stochastic variables. Just like in [2], we say that a job meets the deadline (before we have a realization of the schedule) if the probability that it is completed by the deadline is at least equal to some desired probability; we compute the probability of a job meeting this deadline in a given sequence either analytically or by simulation. A report is to be expected in the near future.

Furthermore, in [147] an overview is presented of the research done in multicriteria scheduling. In [206] we discuss several lower bounds and their efficient computation of the problem of minimizing the maximum lateness on a set of parallel identical machines subject to release dates. Since a bound on the lateness of a job implies a deadline on the completion time of this job, these bounds can be used to decide a feasibility problem that comes up as a subproblem in the no-wait job shop problem.

We have developed several metaheuristic algorithms for practical problems. In [33] we studied the edge-weighted k-cardinality tree problem, which is defined as the problem of finding in a graph the connected component consisting of k edges with minimum weight. This problem arises in applications from facility layout and telecommunications. We present three different metaheuristics for it, based on Tabu Search, Evolutionary Computation, and Ant Colony Optimization. Since the current benchmark instances are too easy, we have developed a new set of benchmark instances on which we have compared our approaches. It turns out that the choice of which method to use depends on the value of k.

In [32] we study the problem of finding the maximum number of edge-disjoint $s_i \to t_i$ paths in a given graph; this problem comes from telecommunications. We present two heuristics for this problem: a multi-start greedy algorithm and an Ant Colony Optimization algorithm. These are tested on a new set of benchmark instances, which come from real-life telecommunication networks. Especially the ACO approach is well suited for solving problems from telecommunications, since it uses only local information for building paths, which is especially important in a dynamic environment. In [91] we studied the (k, r)-center problem in planar and map graphs, which is defined as the question: given k and r, is it possible to choose k centers such that each point is within distance r of at least one center? This problem comes up for instance in the context of facility location, where emergency service facilities have to be installed. We solve this problem by applying a dynamic programming algorithm based on branchwidth, which leads to an algorithm with complexity $f(k, r)n^{O(1)}$, where $f(k, r) = 2^{O(r \log r)\sqrt{k}}$.

Methodological research

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

We have continued our effort on the integration of Constraint Programming and Operations Research. Our findings show that a combined CP and OR approach can increase efficiency, stability and robustness compared to approaches based only on either CP or OR. For the Automatic Recording Problem we use CP-based Lagrangian relaxation ([199]). When a given problem is composed of different substructures for which efficient domain filtering techniques are known, CP-based Lagrangian relaxation allows us to reformulate the problem such that these filtering techniques can be efficiently applied. We propose domain filtering techniques for the maximum clique problem and we present a model in which domain filtering from CP and bounding techniques from OR can be compared ([113]). We give a taxonomy of OR bounds and show that domain filtering outdoes most of these bounds. Using two different branch-and-bound methods from literature we demonstrate the efficiency of domain filtering for the maximum clique problem. In [114] we develop cost-based filtering methods for constrained knapsack problems (CKP). These techniques can be used as basic building blocks when modeling real-life discrete optimization problems. CKPs evolve e.g. when Constraint Based Column Generation is applied to appropriate optimization problems. We develop new linear-time reduction algorithms for Knapsack Problems that are used as propagation routines for the CPK.

In [89] we studied the problem of constructing a good solution for the Traveling Salesman Problem. It is well-known that, if the triangle inequality holds, then the double-tree heuristic finds a solution with worst-case bound 2, independent of which short-cuts are used to construct the tour. We present an algorithm that finds the best set of short-cuts in $O(2^d n^2)$ time and $O(4^d n)$ space, where d is the maximum node degree in the tree, which is smaller than or equal to 6 for the Euclidean TSP; this improves on the time and space requirement of the best algorithm so far. We show that the worst-case performance of this variant of the double-tree heuristic is at least equal to 1.622, and that it outperforms all other construction heuristics when applied to the instances from the TSPLIB database.

WP 4: Generic Methods

Participants: Aarhus, Barcelona, Cologne, INRIA, MPI, Paderborn, CTI, Rome, Utrecht, Warwick (all sites)
Coordinating site: Warwick
Work Package Leader: Mike Paterson
Number of technical reports: 136, of which 72 also appear in other work packages

We summarise below some of the principal areas represented by the research under WP4 during this period. The work is presented in a thematic grouping rather than a strict hierarchical structure (which would be too rigid—a separation into e.g. "Randomization" and "Approximation" would leave no place for a randomized approximation algorithm.)

Smoothed analysis

This is a new area where several ALCOM sites are cooperating in world-class research. Smoothed analysis was first proposed by Spielman and Teng to explain the behaviour of algorithms that work well in practice while performing poorly from a worst-case point of view. The idea is to avoid the sharp peaks of worst cases by some local averaging.

Smoothed competitive analysis

We introduced the notion of *smoothed competitive analysis* of online algorithms in [22]. We successfully applied this to analyze the *Multi-Level Feedback* (MLF) algorithm, which attempts to

minimize the total flow time on a sequence of jobs released over time, when the processing times of jobs are not known until completion. Although MLF performs extremely well in practice its worst-case performance guarantee is exponential. We show that the smoothed competitive ratio of MLF decreases exponentially with the amount of random noise that is added, and give a tight bound for this. A consequence of our analysis is the first average-case analysis of MLF. We show that MLF has a constant expected competitive ratio for several distributions.

In [194], we consider *metrical task systems*, in which an online algorithm resides in a graph of n nodes and may move in this graph at a cost equal to the travel distance. The algorithm has to service a sequence of *tasks* that arrive online; each task specifies for each node a *request cost* that is incurred if the algorithm services the task in this particular node. The objective is to minimize the total request cost plus the total travel cost. Several important online problems can be modelled as metrical task systems, including, for example, the paging problem, the static list-accessing problem and the k-server problem. Borodin, Linial and Saks presented a deterministic *work function algorithm* (WFA) which has a tight competitive ratio of $2n-1$. We give a *smoothed competitive analysis* of WFA, showing that the smoothed competitive ratio of WFA is much better than its worst-case competitive ratio suggests, and that it depends on several topological parameters of the underlying graph.

Knapsack problems

The optimal solution to any given knapsack problem is a Pareto-optimal knapsack filling, i.e., a solution that cannot be improved in weight and profit simultaneously. The knapsack problem can be solved in linear time with respect to the number of its Pareto-optimal fillings. In [26], we investigate the expected number of Pareto-optimal solutions for randomly chosen knapsack instances. This results in the first average-case analysis proving an expected polynomial running time for an exact algorithm for the 0/1 knapsack problem. The random-input model underlying our analysis is very general. We can assume adversarial weights and randomly drawn profits (or vice versa), and can even handle different probability distributions for the profits of different items. This feature enables us to provide a smoothed analysis of the effects of correlations between profits and weights.

Algorithms for moving objects

We propose a new complexity measure for movement of objects [84], the *smoothed motion complexity*. Many applications are based on algorithms dealing with moving objects, but usually data of moving objects is inherently noisy due to measurement errors. Smoothed motion complexity considers this imprecise information and uses *smoothed analysis* to model noisy data. The input is subject to slight random perturbation and the *smoothed complexity* is the worst-case expected complexity over all inputs with respect to the random noise. We illustrate this approach on the problem of maintaining an orthogonal bounding box for a set of n points in $\mathbb{R}^d$ under linear motion. The *motion complexity* is the number of combinatorial changes to the description of the bounding box. We show that the smoothed motion complexity under perturbation with Gaussian normal noise is only polylogarithmic in n.

Randomized algorithms

Selection and sorting

In [176], we analyze simple variants of the Quickselect algorithm, for finding the m^{th} smallest element from n. In the variant *proportional-from-s*, the pivot is chosen at each stage from a sample of size s as the element with the same relative rank as the sought element. Thus if $s = 3$ and we were looking for an element of rank $n/3$, we would choose the smallest element in the sample as our pivot. The analysis shows that proportional-from-3 outperforms median-of-3 for some values of the ratio m/n. The paper analyzes a generalization of proportional-from-3 which we call ν-find, which uses samples of size 3 with interval breakpoints at ν and $1-\nu$, instead of

the 1/3 and 2/3 of proportional-from-3. We find the optimal value of ν. Our results strongly suggest that a suitable implementation of ν-find would be the method of choice in practice. We also consider what happens when $s > 3$, and show that if $s \to \infty$ then proportional-from-s and similar strategies are (almost) optimal.

The short note [173] considers the following common problem: rearrange a given array with n elements so that the first m places contain the m smallest elements in ascending order. We propose a simple variant of quicksort that solves this problem efficiently, and show how it outperforms other common alternatives, such as the partial-sort function in C++'s Standard Template Library.

Primality Testing

Recently, Grantham proposed a randomized primality test and showed that, in the worst case, it had less probability of erring per time spent on the test than the classic Miller-Rabin test. In typical applications, such as the generation of random primes for cryptographic purposes, an average-case analysis is more relevant. Our paper [85] provides such an analysis for an extended version of Grantham's test. This result may lead to faster generation of random primes. This extended algorithm also allows an improved worst-case analysis, using the efficient cubic residuosity test from [86].

Sublinear time methods

Massive data sets are becoming more and more common due to the increasing interconnectivity of modern computer systems. In some applications even linear-time algorithms are too slow to analyze such data sets. Therefore, the long held "gold standard" of linear-time algorithms must be replaced by the search for more efficient "sublinear" algorithms.

One variant of sublinear algorithms is *property testing*, i.e., distinguishing between functions having a predetermined property and functions being *far* from that property. In [82] we show that if the problem of testing a property can be reduced to an *abstract combinatorial program* of small dimension, then the property has an efficient tester. We present efficient property testing algorithms for *geometric clustering*, for the *reversal distance*, and for *graph and hypergraph colouring*. Our framework allows us to analyze all our testers in a unified way and the complexity bounds either match or improve the previously known bounds. We believe that our framework will help to better understand the structure of efficiently testable properties.

In [81] we consider the problem of computing the weight of a Euclidean minimum spanning tree for a set of n points in $\mathbb{R}^d$. We focus on the situation when the input point set is supported by certain basic (and commonly used) geometric data structures that can provide efficient access to the input in a structured way. We present an algorithm that estimates with high probability the weight of a Euclidean minimum spanning tree of a set of points to within $1 + \epsilon$ using only $\widetilde{O}(\sqrt{n}\,\mathrm{poly}(1/\epsilon))$ queries for constant d.

Analytic combinatorics and random structures

Quantifying properties of large random discrete structures is essential in predicting the practical complexity of algorithms. Indeed, over recent years, there has been a gradual shift of emphasis from worst-case scenarios that are often *not* representative of what is observed in actual programming practice to average-case analysis. This trend is reinforced by the spectacular successes of probabilistic algorithms throughout computer science applications.

The report [120] presents in some 130 pages a unified setting for the asymptotic theory of random structures, which is based on generating functions and their complex-analytic properties. This approach is developed in the booklet [78]. It is applied to a wide class of cost functionals of divide-and-conquer algorithms in [117], as well as to urn models of use in determining the expected complexity of balanced trees and B-trees [103]. The problem of connectivity in random graphs is similarly treated in [119]. Patterns in random texts are analysed along such lines in the

reports [121, 191], which find applications in the area of data mining, intrusion detection, and the interpretation of genomic sequences.

A notable new foundational result is [14] where a large class of cost functionals of GCD algorithms see their distributional properties very precisely quantified. This is also of practical relevance, as Euclidean algorithms lie at the heart of many applications involving exact arithmetic, from symbolic manipulation systems to computational geometry and cryptography.

Online algorithms

An important issue when working with online algorithms is how to evaluate their performance. The standard measure for the quality of online algorithms is their *competitive ratio*, which is the worst case ratio of the performance of the online algorithm and the performance of an optimal off-line algorithm. Although the competitive ratio has been widely used for analysis, it has in many cases not been able to differentiate satisfactorily between different online algorithms.

When comparing different online algorithms, a direct comparison seems preferable to an intermediate comparison with the optimal off-line algorithm. In [43], we define a new quality measure for online algorithms, the *relative worst order ratio*, which combines the desirable properties of the max/max ratio proposed by Ben-David and Borodin and the random-order ratio proposed by Kenyon.

In [43] and [44], we analyze several bin-packing problems using this relative worst order ratio, and give some new separations between online algorithms. In [45], we apply the relative worst order ratio to the paging problem. Most "reasonable" deterministic algorithms, including LRU (Least Recently Used) and FWF (Flush-When-Full), have a competitive ratio equal to the size of the fast memory, and this ratio is best possible for deterministic algorithms. We design a retrospective variant of LRU and prove that according to the relative worst order ratio this algorithm is better than LRU, which again is better than FWF. We also prove that, in contrast to the competitive ratio, the relative worst order ratio does reflect that look-ahead can be useful.

In [190] we introduce the online scheduling problem for sorting buffers and present deterministic scheduling strategies for this problem. We are given a service station and a sorting buffer. An input sequence of items characterized by a specific attribute has to be processed by the service station which benefits from consecutive items with the same attribute value. A sorting buffer is a random-access buffer with storage capacity for k items, which can be used to rearrange the input sequence. The goal is to minimize the cost of the service station, i.e., the number of maximal subsequences in its sequence of items containing the same attribute value. The problem is motivated by applications in computer science and economics. We analyze our strategies in the competitive model and show that several standard strategies are unsuitable for sorting buffers. Our main result is a deterministic strategy with a competitive ratio of $O(\log^2 k)$.

In [21] we investigate the quality of service (QoS) achievable by online scheduling algorithms which are *semi-clairvoyant* (i.e., only use approximate knowledge of the initial processing time of each job). We give a semi-clairvoyant algorithm that is $O(1)$-competitive with respect to average flow time. Thus, in some sense, the QoS achievable by semi-clairvoyant algorithms is competitive with clairvoyant algorithms.

Data structures and algorithms

Suffix trees

Suffix trees and suffix arrays are widely used and largely interchangeable index structures on strings and sequences. Practitioners prefer suffix arrays due to their simplicity and space efficiency, while theoreticians use suffix trees partly due to linear-time construction algorithms. In [152], we close this gap between theory and practice by giving a simple linear-time construction algorithm for suffix arrays. We also adapt the algorithm for several parallel and external models of computation. The algorithms for BSP and EREW-PRAM models are asymptotically faster than all previous suffix tree or array algorithms. Based on partly the same ideas, we give in [56] a $O(n \log n)$ time

suffix array construction algorithm that is comparison-based, space-efficient and fast in practice, making it one of the best practical algorithms.

Graph algorithms

A deterministic algorithm for maintaining a t-spanner of a general weighted undirected graph under a sequence of edge insertions and deletions is given in [11].

For the minimum feedback-arc set problem and the minimum feedback-vertex set problem, we give simple combinatorial algorithms for these problems that achieve an approximation ratio bounded by the length of a longest simple cycle of the digraph [96].

Memory management

The binary buddy system is used in many modern operating systems for managing dynamic allocation of small memory blocks whose sizes are powers of two. The standard buddy system allocates and deallocates blocks in $\Theta(\log n)$ time in the worst case and also in the amortized sense, where n is the size of the memory. In [54], we present schemes which improve the running time to $O(1)$ in the amortized and in the worst-case sense.

Algorithms in bioinformatics

Building evolutionary trees is one of the core problems of bioinformatics. The refined Buneman method aims at constructing trees where all edges are well supported by the input distance matrix. The previously best algorithm for computing the refined Buneman tree from a given distance measure had a running time of $O(n^5)$ and a space consumption of $O(n^4)$ which is impractical for realistically sized data sets. In [52], we present an algorithm with time $O(n^3)$ and space $O(n^2)$. The improved complexity makes it possible to explore the virtues of refined Buneman trees on practical data and compare its performance against other methods such as neighbour-joining.

Succinct data structures

The paper [134] gives lower bounds on the tradeoff between the redundancy and the query time for succinct data structures for natural storage and retrieval problems. In particular, the paper gives the first super-polylogarithmic lower bounds on the query time for an explicit static data structure problem, assuming a non-trivial upper bound on the redundancy. The paper received the EATCS best-paper award for Track A papers at ICALP'03.

Complexity

Complexity classes

In [146], the relationship between depth-restricted and width-restricted computation was studied, in the setting of low-level circuit complexity. A notion of width-restricted computation was defined by imposing a geometric constraint of cylindricality on circuits and branching programs. The power of this notion of computation was placed within the hierarchy of complexity classes defined by constant depth circuits between $\mathbf{AC}^0$ and $\mathbf{ACC}^0$. The latter in many ways is the current frontier for circuit lower bounds. In [145], this work was continued by studying a more relaxed notion of geometric restriction. This resulted in the first alternative characterisation of $\mathbf{ACC}^0$ by a circuit model, namely in terms of constant-width polynomial-size planar circuits.

Resolution-width is a key concept in the study of the proof complexity of resolution. It has been used to unify all previously known lower bounds, and also for designing proof-search algorithms of various complexities. It is shown in [9], that resolution-width can be characterized in terms of pebble games. Using this new connection, we solve a well-known open problem in the complexity of resolution, namely, that resolution-width always bounds resolution-space from below. It is hoped that the connection between the two fields will find applications in both directions.

Many problems in interval arithmetic lead to a logical formulation in the framework of first-order logic over the reals, allowing solutions by quantifier elimination. However, the general task of quantifier elimination is known to be NP-hard. In [178], we study a number of such problems, including semi-infinite optimization problems and linear interval equations. We obtain a refined analysis giving precise complexity results in real number models as well as new results in the Turing model.

Completeness in differential approximability classes has been studied in [10]. In differential approximation, the quality of an approximation algorithm is the measure of both how far the solution computed is from a worst one and how close it is to an optimal one. We have defined natural approximation-preserving reductions and proved completeness results for the class of NP optimization problems (class NPO), as well as for DAPX, the differential counterpart of the class APX.

Constraint satisfaction problems

Width notions have been studied in the framework of constraint satisfaction problems in order to exhibit subclasses of problems which allow backtrack-free solution algorithms. In [177], we study a generalization of two such notions introduced by Freuder and Faltings to problems with algebraic constraints. We show that an analogous notion of k-consistency guarantees backtrack-free solution algorithms for tree-structured problems, but that already for binary constraints and a tree-structured constraint graph there arise unavoidable complexity problems in achieving k-consistency. We propose a new width notion which in certain situations even allows us global constraints without yielding a complexity explosion—something not true in the setting above. The approach can be used to find new classes of efficiently solvable 0-1-programming problems.

Temporal reasoning has always been an essential part of artificial intelligence research. Representation of the temporal behaviour of a system via constraints on the time intervals associated with events in this system allows one to retrieve information about possible temporal scenarios for the system. The most popular formalism in temporal constraint reasoning is Allen's interval algebra which concerns qualitative binary relations between intervals (such as "precedes or overlaps"). Deciding whether a given collection of binary qualitative constraints on a set of interval variables is satisfiable is NP-complete in general, and, during the last ten years, many researchers have studied restrictions of this general problem that give rise to tractable problems. A complete classification of complexity for problems with restricted types of allowed temporal constraints was thought to be very difficult to obtain. However this task was accomplished in [162] (in JACM), and further extended in [161] (to appear in SIAM J. on Discrete Mathematics) to the case where additional constraints on interval lengths or time points are available.

Algorithms for NP-hard problems

XSAT is a variant of SAT that is also NP-complete. In 1981, Monien et al. published an algorithm solving XSAT in time $O(2^{0.2441n})$. This is the previously best result for XSAT, but in recent years several algorithms solving X3SAT[5] have been published, the previously best solving it in time $O(2^{0.1626n})$. In [61], we present new improved algorithms for XSAT and X3SAT, running in time $O(2^{0.2325n})$ and $O(2^{0.1379n})$, respectively.

The "exponential" blow-up in size that may occur when converting between CNF and DNF representation of Boolean functions is quantified fairly precisely in [182] as a function of the number of variables and the number of clauses in the CNF. Such bounds are motivated by their fundamental nature and relevance for understanding satisfiability algorithms.

The problem of finding the chromatic number of a graph is a fundamental problem in computer science. The problem is NP-complete but is also provably very hard to approximate. In [60], we give an algorithm for finding the chromatic number, running in time $O(2.4023^n)$. This improves the previously best algorithm by Eppstein running in time $O(2.4150^n)$.

[5] X3SAT is the variant of XSAT in which each clauses contains at most three literals.

Network routing and game theory

Randomized strategies in Nash equilibria

We have studied *randomized strategies* for several games that model selfish routing over non-cooperative networks, with the goal of understanding which ones maximize measures of global performance such as *Maximum Social Cost* and *Quadratic Social Cost* [133, 171, 172]. Our study led us to a tantalizing combinatorial conjecture about Nash equilibria (for two different models of selfish routing). We have termed this the *Fully Mixed Nash Equilibrium Conjecture,* henceforth abbreviated as the **FMNE** *Conjecture.* Formally, the **FMNE** Conjecture states that the *fully mixed Nash equilibrium,* when it exists, achieves the Worst Social Cost. Despite the simplicity and intuitive appeal of the **FMNE** Conjecture, we have been unable to prove it in full generality, but only in several special cases.

Approximation algorithms for network games

We give several approximation algorithms and inapproximability results in [133] for computing *best* and *worst* Nash equilibria, ones that minimize or maximize, respectively, some measure of global performance. (These problems have been shown to be NP-complete.) Some of our approximation schemes are based on the powerful idea of *Nashification* – gradually turn a good approximation of the best Nash equilibrium to an "arbitrarily good" Nash equilibrium.

In [167] we present novel randomized cost-sharing methods for connected facility-location games that are cross monotonic, competitive, and recover a constant fraction of the cost of the constructed solution. Our cost-sharing methods are simple, and achieve attractive approximation ratios.

Randomized bufferless routing algorithms

We present efficient, randomized algorithms for two kinds of *bufferless routing algorithms,* namely *hot-potato routing algorithms* and *direct routing algorithms.* A hot-potato routing algorithm is *greedy* if packets are advanced from their sources towards their destinations whenever possible. We consider greedy, hot-potato routing algorithms for the specific case of a *tree* with n nodes. We give a randomized algorithm, which has a routing time which is within a factor $O(\lg^2 n)$ of optimality with high probability [58].

Direct routing is a special case of bufferless routing in which packets must be routed *without conflicts* along specific paths. We show in [57] a *generic* randomized greedy direct routing algorithm with routing time that is worst-case optimal. We have shown that variants of the greedy algorithm achieve *optimal* or *near-optimal* (within logarithmic factors) routing times for many interesting routing problems on commonly used network topologies such as trees, meshes, hypercubes and butterflies.

WP 5: Experimental Algorithmics

Participants: Aarhus, Barcelona, Cologne, INRIA, MPI, Paderborn, CTI, Rome
Coordinating site: MPI
Work Package Leader: Peter Sanders
Number of technical reports: 40, of which 31 also appear in other work packages

The large number of reports in WP 5 that are also appearing in at least one other work package underlines the fact that algorithm engineering has now become an integral component of algorithmic research.

In the following, we describe selected areas where algorithm engineering played a particularly big role in the third period of the project. Note that some engineering aspects also appear in the descriptions of other work packages.

MPII organized a summer school focused on algorithm library design that was partially funded with ALCOM money[6]. Four $2 \times 1/2$-day mini-courses were given:

- Robert Bixby and Ed Rothberg (Ilog): The CPLEX Library for Linear, Mixed-Integer, and Quadratic Programming.
- Dan Halperin (Tel Aviv): Arrangements and Their Applications.
- Bertrand Meyer (ETH, Eiffel Software): Principles of Library Design, from Reuse and Contracts to Proofs: the Eiffel Experience.
- Stefan Näher: Design and Implementation of Efficient Data Types for Static Graphs.

Data Structures

Very often, the difference between a naive and an advanced algorithm is the use of a sophisticated data structure. However, these data structures often look too complicated, too theoretical, or are simply unknown to application designers. Since a single data structure could impact many different applications we therefore study the performance of advanced data structures.

One of the most versatile data structures are representations of dictionaries by search trees that allow fast insertion, deletion, and search. Since the 1970s, data structures have been known that support all these operations in time $O(\log \log C)$ for integers in the range $0..C$. However, we knew of only a single implementation study (part of a master thesis at MPII in 1992) which concluded that these data structures cannot practically compete with simpler comparison based alternatives. Since this is an apparent gap between theory and practice, we have studied this problem more closely [92]. It turns out that a tuned implementation reverses the picture. For 32-bit keys and large sets, the tuned integer data structure searches about twice as fast as the best comparison based implementations.

At the cost of less versatile searching, hash tables allow an even faster implementation of dictionaries. One gap between theory and practice here was that dynamic dictionaries that allow search in worst case constant time (e.g., for real-time applications) implied a waste of space by a considerable constant factor. In a cooperation between Aarhus, MPII, and Patras we have developed a simple generalization of cuckoo hashing that solves this problem [128]. For example, raising the number of memory probes from two to three reduces the waste of space from over 50 % to less than 10 %.

In the following section we report on a concrete application of geometric data structures on an important problem in bioinformatics.

Evolutionary Trees

Building evolutionary trees is one of the core problems bioinformatics. The neighbor joining method by Saitou and Nei is one of the most widely used methods for building such trees from information about the pairwise distances between a set of species. The standard neighbor joining algorithm has $O(n^3)$ worst case running time, which makes it impractical to use for building trees for large families of e.g. vira. In [51], we present a different way of implementing the neighbor joining algorithm. It constructs the same evolutionary trees as the standard algorithm but is faster on the average. We empirically evaluate its performance on distance matrices obtained from the Pfam collection of alignments by comparing its running time with that of the QuickTree tool—a highly tuned and widely used implementation of the standard neighbor joining method. The experiments confirm that our implementation yields a significant speed-up over QuickTree.

[6] `http://www.mpi-sb.mpg.de/~adfocs/`

Cache-Oblivious Algorithms

Cache-obliviousness, described in more detail in the section on Work Package 1, is a new and theoretically attractive way of dealing with the multi-level memory hierarchies found on modern computers. However, empirical investigations of its merits have been scarce so far. In [53], we perform such an investigation in the area of comparison based sorting.

We explore a number of implementation issues and parameter choices for the cache-oblivious sorting algorithm Lazy Funnelsort, and settle the best choices through experiments. We then compare the final algorithm with Quicksort, generally acknowledged to be the fastest all-round comparison based sorting algorithm. The end result is a carefully implemented cache-oblivious sorting algorithm, which our experiments show can be faster than the best Quicksort implementation we can find, already for input sizes well within the limits of RAM. On disk the difference is even more pronounced regarding Quicksort, whereas it is slower than a careful implementation of the (cache-aware) multiway Mergesort algorithm.

This demonstrates that the overhead involved in being cache-oblivious can be small enough for the theoretically nice memory usage properties to transfer into practical advantages over algorithms developed for the classic RAM model.

Uniform Random Generation of Combinatorial Objects

Work otherwise carried out in WP4 has led to a very precise characterization of what to expect regarding the number of occurrences of patterns of a possibly complex shape in random texts obeying a variety of probabilistic models (e.g., [121]). This makes it possible to confront analysis against real data and detect meaningful signals in large data sets. This is the purpose of [168] where extensive data on genomic sequences of plants are tested in this light.

The report [103] proposes a new approach to the structure of random objects based on the idea of Boltzmann models from statistical physics. The algorithms are almost invariably of a subquadratic complexity, often linear or quasi-linear, and they can be automatically generated from specifications. Prototype implementations of [103] demonstrate that they can be put to use in order to simulate models of sizes up to 100,000, in situations where only toy simulations of size 100 at most where previously possible. Generators can be constructed for varieties of constrained sequences, trees, graphs, and permutations, and applications are envisioned in the area of random testing.

A major new result is [105] (also mentioned in WP1), by which one can estimate cardinalities of extremely large data sets using a single pass and one or two kilobytes of memory only. The effectiveness of the algorithm has been extensively tested on a whole range of data up to a gigabyte including server and router traces, and it fully confirms the soundness of the algorithm on real-life data.

Graph partitioning

We have continued to improve our graph bisection heuristics. To better compare our solutions to those obtained using other available packages like Metis and Jostle, a new node permutation based evaluation scheme was developed [196]. Based on this new scheme, we were able to locate and enhance critical sections in our helpful-set implementation and thus lower the edge-cuts computed by our graph partitioning library PARTY by up to 20%, while its run-time only increases by a small amount. In [195], we have shown that our helpful-set heuristic is also capable of computing partitions with a very high imbalance factor. Comparing it to other state-of-the art libraries, we obtain much better results here.

Another approach to partition a geometric graph is applying space-filling curves that order all vertices linearly. The domains are then defined according to the computed ordering. This approach is very fast and requires only very little memory, but since the edges of the graph are ignored the resulting edge-cut is higher. In [197] we have considered several types of space-filling curves and compared the quality of their induced partitionings to the one computed by Metis. One

can see that applying the latter results in around 30% to 50% less cut edges for a small number of partitions while for larger numbers this advantage decreases.

The Concept of Treewidth

The graph theoretic notion of treewidth has played a major role in several theoretical investigations in the past years. Its role in practical implementations however has been small so far. We aimed at developing methods that make this notion also well usable in practical application contexts [36, 37, 107]. One issue to be solved is to obtain good methods to find tree decompositions of given networks with optimal or close to optimal treewidth. We developed different methods for preprocessing these networks, either based on reduction rules or on methods to split the graphs in smaller parts with help of safe separators. The methods were shown to reduce the networks to smaller but still equivalent instances. Tests on networks derived from applications like probabilistic networks (a technology underlying several modern decision support systems) and frequency assignment have shown that the methods are usually very effective.

Dynamic Graph Algorithms

In [94] we have presented the results of an extensive computational study on dynamic algorithms for all pairs shortest path problems. We have described our implementations of the recent dynamic algorithms of King (FOCS 1999) and of Demetrescu and Italiano (ALCOMFT-TR-02-92), and compared them to the dynamic algorithm of Ramalingam and Reps (Journal of Algorithms 21, 1996) and to static algorithms on random, real-world and hard instances. Our experimental data suggest that some of the dynamic algorithms and their algorithmic techniques can be really of practical value in many situations.

Algorithm Visualization, Monitoring, and Debugging

As a follow-up of a previous work done during the ALCOM-FT project (ALCOMFT-TR-01-177), we have developed Leonardo Web, a collection of tools for building animated presentations that can be useful for investigating and disseminating algorithmic concepts [40]. Presentations can be created via the combined use of a visual editor and a Java library. The library allows it to generate animations in a batch fashion directly from Java code according to an imperative specification style. Batch-generated animations can then be refined and customized using the editor. Presentations can be finally viewed with a simple Java player, which ships both as a stand-alone application for off-line deployment and as a Java applet embedded in a Web page. The player supports step-by-step and continuous execution, reversibility, speed selection, and smooth animation.

The intrinsically dynamic nature of running programs, opposed to static implementation code, makes programming, debugging, and performance monitoring difficult and time-consuming. Development environments typically provide source-level debuggers that help see things in action, clarifying how the code is interpreted by a computer, but do not help much with portraying the behavior of complex programs at a more abstract level neither do provide advanced tools for fault localization. Devising powerful methodologies and flexible tools for assuring software reliability still remains a very difficult problem. In [40] we describe a cross-platform virtual machine that provides advanced facilities for implementing directors, i.e., reactive systems that monitor the run-time environment and react to the emitted events. Typical examples of directors are debuggers and tools for program analysis and software visualization.

Graph coloring algorithms

In [5], we design and implement two algorithms for coloring the square of planar graphs (SQPG). We call these algorithms MDsatur and RC. We have also implemented and experimentally evaluated the performance of most of the known approximation coloring algorithms for SQPG. We

compare the quality of the colorings achieved by these algorithms, with the colorings obtained by our algorithms and with the results obtained from two well-known greedy coloring heuristics. The heuristics are mainly used for comparison reasons and, unexpectedly, give very good results. Our algorithm MDsatur outperforms the known algorithms as shown by the extensive experiments we have carried out. The planar graph instances whose squares are used in our experiments are "non-extremal" graphs obtained by LEDA and hard colourable graph instances that we construct.

Distributed Algorithms

We have designed, implemented and evaluated (both by analysis and large scale simulation) several new algorithms for important problems in non-conventional networks (mostly ad-hoc mobile and wireless sensor networks).

For the simulations, we have chosen to use LEDA and have written specific C++ code extensions in each case. The choice of such a lightweight environment allows to simulate very large networks of thousands of users — this is very important in order to investigate the scalability of our protocols. On the contrary, other (more heavy) simulation environments (like ns-2), are more detailed in terms of the network specifications they capture, however the size of inputs studied is rather limited.

Implementing our protocols and evaluating them has been valuable in the following respects: a) it allowed us to precisely evaluate the (difficult to analyze) particular impact of several parameters and the exact evolution of some phenomena b) in some cases it revealed unexpected behavior of designed protocols, that further guided our analysis and design efforts c) we got some preliminary hints about a possible implementation of protocols in real (current technology) devices and networks d) we could comparatively evaluate different protocols.

Wireless Sensor Networks: We have focused on the problem of efficient propagation of data to a fixed control center, achieved by the distributed co-operation of the sensors in the network. We have designed several protocols for this problem: the Local Target Protocol (LTP, performing a local optimization in selecting the next hop sensor, [70]), the Probabilistic Forwarding Protocol (PFR, that probabilistically favors close to optimal transmissions, [71]), the Sleep-Awake Protocol (S-W, that aggressively puts sensors to periodically switch-off in order to save energy, [75]), the Variable Transmission Range Protocol (VTRP, that may increase transmission range to bypass obstacles, critical regions or faulty sensors, [7]). We have investigated various performance and fault-tolerance efficiency measures (like the success probability of the protocol, the number of data transmissions, the time needed, the available energy, the number of sensors that become involved in the propagation process, the number of "alive" particles etc.), for various parameters (like the density of sensors in the network area, the number of time steps, the number of events to be reported etc).

Ad-hoc Mobile Networks: In [72] we demonstrate the significant impact of (a) the mobility rate and (b) the user density, on the performance of routing protocols in ad-hoc mobile networks. The main findings are that protocols based on route-maintenance behave well in networks of high user density and low mobility rate, while their performance drops for sparse networks of highly mobile users. On the other hand, protocols that instead use accidental meetings of hosts to forward messages tolerate well (and in fact benefit from) high mobility rates and low densities. In [74] we study systems of anonymous and oblivious robots each of which observes the positions of the other robots and moves to a new position so that eventually the robots form a circle. We propose a new distributed algorithm that, as shown by an extensive and detailed comparative simulation experimental study is very simple and takes considerably less time to execute than algorithm DK in the recent literature.

In a similar line of research, we, computer scientists, are working on the three parts modelling and theoretical investigation, simulation, and finally prototypical realization in conjunction with electrical engineers. This collaboration is extraordinary since we try to combine algorithm

development and mathematical analysis, experimental investigation and, finally, the prototypical implementation. By this the physical effects of the wireless data communication have been modelled and integrated in our simulation environment to obtain simulations very close to reality.

In [208] we presented a new model for communication in MANETs. Instead of omnidirectional transmissions, as assumed in most papers and all existing systems, the members are allowed to submit data in a fixed number of different directions (sector subdivision) and to adjust the transmission power in each sector separately. A simulation environment (Simulation Environment for Ad Hoc Networks, SAHNE) [192] was presented that allows simulation of communication strategies in MANETs that use sector subdivision, and simulation results will be shown where communication paths are selected via hop-minimization or geometric spanner properties. SAHNE is based on C++ and common libraries, which ensures that it can be used on many different platforms. The experiments showed the influence of different parameters in realistic scenarios, and using geometric routing seemed to be better than using hop-minimization.

DAP: A Generic Platform for the Simulation of Distributed Algorithms: In [73] we describe a generic and homogeneous simulation environment aiming at the implementation, simulation, and testing of distributed algorithms for wired and wireless networks. We present its architecture, the most important design decisions, and discuss its distinct features and functionalities. DAP allows the algorithm designer to implement a distributed protocol by creating his own customized environment, and programming in a standard programming language in a style very similar to that of a real-world application. DAP provides a graphical user interface that allows the designer to monitor and control the execution of simulations, visualize algorithms, as well as gather statistics and other information for their experimental analysis and testing, respectively.

WP 6: Project Management, Dissemination, Evaluation

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

Project Management

The Consortium Board, consisting of the scientific leaders of each of the 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. The Coordinator has also been the point of contact between the EC and the consortium, 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 Leaders, who has also conducted the reporting on the work done.

The ALCOM-FT Consortium Board met on July 5, 2002 in Warwick. The Board found the project to be in overall good shape—the productivity in terms of scientific papers was high, and deliverables were being delivered on schedule, except for a few delayed due to hiring problems. In total, EUR 1.270.000 out of the project budget of EUR 2.160.000 had been spent by the end of the second project year. This was a bit lower than planned, but in good correspondence with the described hiring problems. In particular, coordination as well as sites MPI and Utrecht had under-spent, site Cologne had overspent, and the newly joined site Cyprus had not spent yet. The remaining sites had all spent approximately the planned 2/3 of the budget. A minor transfer in

the budget of funds from coordination and site MPI to site Cologne was agreed on, in order to ensure the completion of the deliverables produced at site Cologne.

The possibility of an extension of the project with half a year (no additional funding) was discussed. The Board decided to apply for this, to ensure completion of the delayed deliverables. This extension was later granted by the Commission.

During the third project period, the hiring problems have been solved, and all deliverables have been finished by the end of the extension. An additional benefit of the extension have been to increase the time available for site Cyprus, which joined the consortium March 1, 2002, to integrate into the consortium. The addition of site Cyprus has strengthened the ALCOM-FT project in the areas of distributed computing and networks, and has in particular benefited the work in Work Packages 2 and 4. Site Cyprus has been well integrated in the consortium, as is witnessed by the amount of cooperation between site Cyprus and other ALCOM sites (see page 38).

The second annual ALCOM-FT review meeting and workshop was held in Warwick, July 5–6, 2002. The review meeting featured a presentation of the work done in each of the work packages of the project. At the preceding scientific workshop, each site presented details of selected work from the second year of the project.

Dissemination

The principal means of dissemination for the ALCOM-FT project is scientific reports describing the research done. A total of 206 reports have been produced during the third period of the project. They have all appeared in the ALCOM-FT Technical Report Series. All deliverables of the project, including the ALCOM-FT Technical Report Series, are available through the ALCOM-FT website at `http://www.brics.dk/ALCOM-FT`.

During the third project period, four summer schools were organized by consortium members with partial funding from ALCOM-FT. These were: *Models and Algorithms for the Web (MAW 02)*, held in Udine, June 17–22, 2002, *3nd Max-Planck Advanced Course on the Foundations of Computer Science (ADFOCS 2002)*, held in Saarbrücken, September 9–13, 2002, *Workshop on Parallelism in Algorithms and Architectures*, held in Paderborn, March 10, 2003, and *4th Max-Planck Advanced Course on the Foundations of Computer Science (ADFOCS 2003)*, held in Saarbrücken, September 8–12, 2003.

In September 2002, site Rome organized the main European event in algorithmics, the federated conference ALGO 2002 in which ESA 2002, APPROX 2002 and WABI 2002 were colocated. In July 2003, site Aarhus organized the major complexity conference CCC 2003. Additionally, many smaller conferences and workshops with subjects relevant for the project were arranged by various ALCOM-FT sites. These are reported in Part II of this report (starting page 32).

A website (deliverable D14) with *PR-pages for algorithmics* has been set up, and can be found at `www.mpi-sb.mpg.de/~spyros/pr-www/`. The purpose of this site is to create a virtual meeting point between the communities of computer scientists and practitioners. This meeting point should aid the flow of information and know-how between the European sites of computer scientists and the industry. More specifically, the site reports success stories proving by example the applicability of rigorous theoretical expertise in practice, it provides a repository of problems and techniques helping people in industry asking the right questions about their real-life problems and contact the right persons in CS community, and it provides state-of-art know-how through plenary talks given by computer scientists on main topics of research in computer science.

We also continue to disseminate algorithmic awareness through talks given by partners to audiences in industry. For example, site Paderborn has expanded their list of talks in their 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 `www.uni-paderborn.de/cs/kooperation/iti.html`).

The consortium continues to maintain a high profile in PhD education, illustrated by the fact that 18 persons obtained an ALCOM-FT related PhD-degree during the last project period.

On the negative side, the algorithmic competition (see `www.mpi-sb.mpg.de/~schaefer/MLLB/index.html`) set up in the last project period in order to increase the awareness of algorithmics

in research communities outside ALCOM-FT have been canceled due to lack of entries submitted. Apparently, this venue of dissemination is not efficient, at least not in the form attempted here.

Evaluation

The ALCOM-FT project has finished the last period in overall excellent shape. All deliverables have been completed. Scientific productivity has been very high, resulting in a total of 206 reports. The international standing of ALCOM-FT is illustrated by the fact that 146 of the 206 reports have already[7] been published or been accepted for publication in the scientific community via conferences and journals, including the following key conferences: ALENEX, CCC, CPM, ESA, EURO-PAR, FCT, FOCS, ICALP, ISAAC, LATIN, MFCS, SPAA, SODA, STACS, STOC, WABI, WADS, and WEA. Cooperation between sites has been good, as can be seen from the reporting from the individual sites. Some deliverables were delayed due to problems with filling programming positions, but these problems were handled by the six month extension.

[7]Based on the reporting from the authors. This is actually a lower bound, since most likely some publications have not been reported.

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 Southern Denmark. The site is the coordinating site of the ALCOM-FT project, and is involved in Work Packages 1, 4, 5, and 6. The group in Aarhus is part of BRICS (Basic Research in Computer Science, Centre and International Research School of the Danish National Research Foundation).

In total, the two groups currently comprise 11 associate/assistant professors and 7 Ph.D. students. Rolf Fagerberg is paid part-time by ALCOM-FT funds, and is responsible for running the day-to-day business of the project. During the third project period, three Ph.D. degrees were obtained based on ALCOM-FT related work: Lene Monrad Favrholdt (thesis *On-Line Problems with Restricted Input*), Morten Nyhave Nielsen (thesis *On Measuring the Quality of On-Line Algorithms*), and Rasmus Pagh (thesis *Hashing, Randomness and Dictionaries*).

Visitors to the site include Riko Jacob (ETH, Zürich), Lukasz Kowalik and Maciej Kurowski (Faculty of Mathematics, Informatics and Mechanics, Warsaw University), Thore Husfeldt (University of Lund), Anna Östlin and Rasmus Pagh (IT University of Copenhagen) John Iacono (Polytechnic University), Lars Arge (Department of Computer Science, Duke University), Seth Pettie (University of Texas at Austin), Martin Dietzfelbinger (Technische Universitaet Ilmenau), Jaikumar Radhakrishnan (Tata Institute), S. Srinivasa Rao (University of Leicester), Erich Novak (Friedrich Schiller Universität Jena), Marco Gori, (University of Siena), Rene Peralta (Yale University), Susanne Albers (Freiburg University), and Gerhard Woeginger (University of Twente). Rasmus Pagh visited site MPI, and Bjarke Skjernaa was visiting scholar at site MPI for six months.

Peter Bro Miltersen served in the program committee of the conferences FSTTCS 2003 and STACS 2003. Gerth S. Brodal served in the program committee of the conferences LATIN 2004, SODA 2004, WADS 2003, EUROMICRO-PDP 2003 (special session on memory hierarchies), and ALENEX 2003. Rolf Fagerberg served in the program committee of ESA 2004. Joan Boyar served in the program committee of WEA 2003.

In July 2002, the site organized and hosted the *EEF Summer School on Massive Data Sets*, in January 2003 the *Tenth Workshop on Applied/Advanced Research in Combinatorial Optimization*, and in July 2003 the 18th *IEEE Conference on Computational Complexity* (CCC 2003).

Members of the site have participated in most of the important algorithm conferences, including ESA, FCS, FOCS, ICALP, ISAAC, SODA, STOC, SWAT, and WABI.

The 23 ALCOM-FT reports [28,43–45,49–54,60,61,85,86,128,129,134,145,146,156,177,178,182] 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. During this second period, the site contributes to the following Work Packages: Work Package 1 with the implementation (and, in some cases, design) of sampling-based efficient data mining algorithms and in the tasks of "Communication primitives;" In Work Package 2: we have done work in the tasks of Communication primitives; Static and dynamic load balanced; In Work Package 4 we have made contributions to the tasks of Randomness models and probabilistic analysis; Graphs and networks; Complexity bounds; and New Data Structures. We also have made some contribution to Work Package 5. In total, the group currently comprise 10 associate/assistant/full professors and 3 Ph.D. students.

Barcelona is actively collaborating in research with INRIA (P. Flajolet) and CTI (P. Spirakis). Other places where there is a collaboration are Cambridge University, University of Bergen, MIT,

McGill University, Waterloo University and University of New Mexico at Sante Fe, Tokyo Institute of Technology. Visitors to the Barcelona site include , Philippe Flajolet (from INRIA), Paul Spirakis (CTI), Cristopher Moore, Martin Loebl, Joerg Flum, Jose Rolim, Yuri Gurevich, Denis Therien. On the other hand, C. Martinez visited INRIA Paris, J. Gabarró visited INRIA-Paris A. Atserias visited University of Edimburgh, Cambridge University and UC Santa Cruz, R. Gavalda visited McGill University, D. Thilikos visited MIT, J. Balcazar visited the Univ. of Illinois at Urbana-Champaign.

A. Atserias gave two expository invited talks, one on the *Workshop on Typical Case Complexity and Phase Transitions*, and the other *Workshop on Implicit Computational Complexity*, both satellite workshops of LICS 2003.

During the present period, 32 Technical Reports were produced within ALCOM-FT [8, 9, 12, 13, 15, 32, 33, 35, 67–69, 83, 90, 91, 97–99, 118, 122–125, 131, 135, 144, 169, 173–176, 188, 202].

Most of this work was presented or accepted for publication in the following meetings and journals: SODA-03, IEEE Conf. on Computational Complexity-03, STACS-03, ICALP-03, ESA-03, MFCS-03, EUROCOMB-03, WG-03, STACS-04, LATIN-04, Euro-Par-03, SODA-04, Discrete Math. and Data Mining-03, Multi relational Data Mining-03, Journal of the ACM, Theoretical Computer Science, Parallel Processing Letters, Information and Computation, IEEE-Transactions on Mobile Networks, Computers & Operations Research and Formal Aspects of Computing.

R. Gavalda was a co-chair of the PC of ALT'03, A. Atserias is on the PC on ICALP-04 and LICS-04, J. Díaz is on the PC of ICALP-04 and MFCS-04.

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. During the third period the group compromises 5 Ph.D. Students. Christoph Buchheim completed his Ph.D. with the thesis on "An Integer Programming Approach to Exact and Fuzzy Symmetry Detection". Matthias Elf was payed full-time by ALCOM-FT funds. The Cologne site had bee involved in finishing deliverables D20, D23, and D29.

Matthias Elf visited the MPI for 3 month. Members of the site attended GD 2002, GD 2003, OR 2002, OR 2003, ISAAC 2002 and the Aussois Workshop on Combinatorial Optimization

Michael Jünger has been elected to the board of the German Society of Operations Research. Two books had been produced at the Cologne site, *Combinatorial Optimization – Eureka You Shrink!* (ISBN 3-540-00580-3) and *Graph Drawing Software* (ISBN 3-540-00881-0). Both had been published at Springer.

The Cologne Site produced the following ALCOM-FT relevant reports [16, 55, 150, 170]. Of these, one has been accepted for and international conference, one has been published in an international journal and one is to appear as a chapter of a book.

INRIA

The activity at site INRIA (Rocquencourt, France) is based on the Algorithms Group, a research project focusing on the design and analysis of algorithms, and the RAP Group whose research focuses on probabilistic methods, networks, and communication. Altogether, he INRIA team of Alcom-FT comprises 5 permanent members (one junior and 4 senior members), 4 PhD students, and 4 associate members (from Universities of Caen and Paris). The research of three of the PhD students (M. Durand, J. Fayolle, and V. Puyhaubert) takes place entirely within the framework of the Alcom Project. The work packages concerned are WP1, WP2, WP4, and WP5.

Research produced in the group continues to be directed primarily towards general methods aiming at precisely quantifying randomness in large discrete structures (WP4). The framework developed is largely 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 130 pages has been produced [120] regarding complex-analytic methods [WP4] and is part of a long haul book project on *Analytic Combinatorics* with R. Sedgewick (Princeton). We have developed an extensive visitor programme with publication of the *Algorithms Seminar* proceedings as a booklet [78] (190 pages) resorting to WP1, WP2, WP4, WP5. New directions in the random generation of combinatorial structures [WP4, WP5] have been developed to the point of becoming operational [103]. A notable new theoretical result [WP4] is the distributional analysis of Euclidean algorithms [14]. A notable practical result [WP1, WP5] is the design of the best known algorithm for estimating cardinalities of data sets with applications in data mining and router activity monitoring [105]. Applications to networks [WP2] are being developed jointly with the Research Division of the French Telecom: tangible progress has been made in modelling very precisely ADSL traffic on IP [27] and a book (*Stochastic Networks and Queues*, by P. Robert, Springer, 398 pages, 2003) has recently appeared in print. We have played a major rôle in the organization of the *International Seminars on Analysis of Algorithms* (Strobl 2002; San Miniato, 2003) and in the setting up of the French network ALÉA (Maresilles, 2003) that are specifically dedicated to random structures and analysis of algorithms [WP4]. Each event gathered more than 70 researchers worldwide.

The INRIA site has produced 16 ALCOM-FT reports

[14, 27, 46, 78, 103–105, 117–121, 143, 168, 191].

Of these 14 have been submitted to international journal and conferences, with acceptance at the ICALP, ESA, ANALCO, INFOCOM conferences and by journals like *Combinatorics, Probability, and Computing*, *El. J. Combinatorics*, *Applied Prob.*

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, 19 researcher with Ph.D. (2 with habilitation), and 17 Ph.D. students. More than half of them work partially on ALCOM-FT. Roman Dementiev and Ulrich Meyer were paid from ALCOM-FT funds.

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 Guy Even, Rasmus Pagh (Aarhus), Bjarke Skjernaa (Aarhus), Sandeep Sen, Paul Spirakis (CTI), and Stefano Leonardi (Rome).

Members of the site have participated in most of the important relevant conferences, including ESA, FOCS, ICALP, SODA, SPAA, STOC, SWAT. The site hosted the third[8] Max-Planck Advanced Course on the Foundations of Computer Science (ADFOCS) in the summer of 2002. Speaker were Timothy Chan (Random Sampling), Michele Mosca (Quantum Computing), Prabhakar Raghavan (Randomized Algorithms and the Probabilistic Method), and Gerhard Wöginger (Approximation Algorithms and Inapproximability). The fourth ADFOCS[9] in the summer of 2003 was on algorithm library design. For more details refer to the section on WP5. Both events were co-funded by ALCOM.

The 22 ALCOM-FT reports [4, 17, 18, 22–26, 56, 92, 93, 108, 127, 128, 130, 148, 152, 153, 160, 165, 167,180,194] 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.

[8] http://www.mpi-sb.mpg.de/conferences/adfocs-02/
[9] http://www.mpi-sb.mpg.de/conferences/adfocs-03/

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 third 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 24 Ph.D. students.

During the third year of ALCOM-FT, Robert Elsässer (*Spectral Methods for load balancing Strategies*), Meinolf Sellmann (*Reduction Techniques in Constraint Programming and Combinatorial Optimization*), Torsten Fahle (*Integrating Concepts from Constraint Programming and Operations Research Algorithms*), Harald Räcke (*Data Management and Routing in General Networks*), Norbert Sensen (*Lower Bounds and exact Algorithms for the graph partitioning problem using multicommodity flows*) Christian Sohler (*Property Testing and Geometry*), and Martin Ziegler (*Zur Berechenbarkeit reeller geometrischer Probleme*) completed their Ph.D.

Visitors to the site include Udo Adamy (ETH Zürich, Switzerland), Sergei Bezrukov (University of Wisconsin-Superior, USA), Christian Breimann (University of Münster, Germany), Artur Czumaj (New Jersey Institute of Technology, USA), Hans Florian Geerdes (ZIB Berlin, Germany), Ralf Klasing (CNRS/INRIA Sophia Antipolis, France), Arie Koster (ZIB Berlin, Germany), Petter Kristiansen (University of Bergen, Norwegen), Marios Mavronicolas (University of Cyprus, Cyprus), Christian Scheideler (Johns Hopkins University, USA), Bhabani Sinha (Advanced Computing and Microelectronics Unit, Calcutta, Indien), Dimitrios Thilikos (University of Barcelona, Spain), Abderezak Touzene (Sultan Qaboos University, Oman), Berthold Vöcking (University of Dortmund, Germany), Philipp Wölfel (University of Dortmund, Germany), and Alex Znamenshchikov (Kharkiv National University, Ukraine).

Within the last year Burkhard Monien visited the ALCOM partners in Cyprus and Rom. Thomas Lücking, Manuel Rode, and Martin Gairing visited the ALCOM partners in Cyprus.

Moreover, within the last year Burkhard Monien has been a member of the program committees of ESA'03, and IPDPS'04 (Program Chair). Friedhelm Meyer auf der Heide was involved in the scientific program committees of SPAA 2003 and he has become managing editor of *Journal of Interconnection Networks*. Rolf Wanka and Christian Schindelhauer have been members of the program committee of STACS'03 and 04, respectively. Harald Räcke received the best paper award and the Machtey award for the best student paper at FOCS'02 for his contribution *Minimizing Congestion in General Networks* (cf. ALCOM-FT technical report [189]). Members of the site Paderborn have participated in most of the important and relevant conferences, including CPAIOR, ESA, EUROPAR, FOCS, ICALP, IPDPS, MFCS, PDPTA, SODA, SPAA, and WADS. 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.

As one can see from the appended list of ALCOM-FT reports we were able to establish during the third period a strong cooperation between our site and the sites of Cyprus and CTI in Patras.

The 34 ALCOM-FT reports [30, 31, 48, 81, 82, 84, 88, 109–116, 133, 141, 142, 154, 155, 171, 172, 181, 189, 190, 192, 193, 195–199, 208, 210], have been authored or co-authored by members of the Paderborn 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 10 faculty members/postdoctoral researchers and 15 PhD students.

During the third year of ALCOM-FT, the PhD Theses of I. Chatzigiannakis (on ad-hoc mobile networks), V. Papadopoulou (on radiocoloring), D. Koukopoulos (on network stability) and

A. Svolos (on parallel image processing) have been completed. Also, P. Kanellopoulos and S. Athanassopoulos have obtained their M.Sc. Degree. Three postdoctoral researchers (S. Kontogiannis, I. Chatzigiannakis and V. Papadopoulou) have now been added to the group. Nine students (C. Raptopoulos, C. Efthimiou, P. Panagopoulou, V. Liagkou, A. Kinalis, G. Mylonas, I. Krommidas, E. Pirga and G. Tsaggouris) have began their PhD/M.Sc.

Members of the site have participated in the Program Committees of major relevant Conferences: P. Spirakis in MFCS 2003, IPDPS 2004, WEA 2004, FCT 2003 and ALGOSENSORS 2004, C. Kaklamanis in ICALP 2004, SPAA 2004, EUROPAR 2003, EUROPAR 2004, IPDPS 2004, CIAC 2003, ADHOCNOW 2004, S. Nikoletseas in PERCOM 2004, WMAN 2004, WEA 2003, C. Zaroliagis in WG 2003, FUN 2004 and ATMOS 2002. S. Nikoletseas co-organizes (with Jose Rolim) the first International Workshop on Algorithmic Aspects of Wireless Sensors Networks (ALGOSENSORS), to be held in conjunction with ICALP 2004.

P. Spirakis has visited C. Papadimitriou in Berkeley and MPI. He was keynote speaker at the Onassis Research Foundation Summer School. S. Nikoletseas was invited to give a tutorial on wireless sensor networks at the 2nd International Technology Symposium (I2TS 2003). C. Zaroliagis was invited to Dagstuhl-Seminars on "Experimental Algorithmics", on "Data Structures", on "Algorithmic Aspects of Large and Complex Networks" and he was distinguished lecturer of the Bertinoro Spring School.

34 ALCOM-FT Technical reports ([5–7,41,42,57,59,62,64–66,70–77,100,106,126,128,133,151, 157–159, 172, 183–187]) have been authored or co-authored by researchers of the site during the third year. Currently, 25 out of these have been accepted by conferences (15 papers, including ESA, STACS, MFCS, ISAAC, European Wireless (EW), EUROPAR, WEA, SIROCCO, the Annual Simulation Symposium (ANSS)), in journals (5 papers, including TOCS, MONET, PPL) and as Chapters in Books of International Circulation (5 Chapters, published by Springer, Kluwer, World Scientific).

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), 2 Ph.D. students and 2 post-docs. Camil Demetrescu is payed as researcher on ALCOM-FT funds; Paolo Terrevoli and Silvana Di Vincenzo are assisting with the coordination of the site.

During the third year of ALCOM-FT, Luigi Laura has completed his Ph.D. with the thesis *Modeling and Mining Complex Information Systems*. Debora Donato has completed the third year of her PhD program and Vincenzo Bonifaci and Luca Alluli have just started their graduate studies. Visitors to our site during this period include Guido Schaefer (from MPI), Tjark Vredeveld, Rene' Sitters and Wilhelm de Paepe (from Eindhoven), Anna Karlin, Amos Fiat, Friedhelm Meyer auf der Heide, Benny Chor, Martin Skutella, Baruch Awerbuch. More recently various foreign postdocs (Alex Kesselman, Jochen Koenemann and Panayotis Tsaparas, Marco Gaertler) have joined the group.

Members of the site have participated in most of the important relevant conferences, including ESA, FOCS, ICALP, SODA, 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 ESA'02 and TCS'02. Stefano Leonardi has been in the Committees of SODA03, SPAA03, ICALP03, APPROX03, STACS04 and TCS04. Giorgio Ausiello has been in the APPROX'02 PC, and Camil Demetrescu in the ESA'02 PC.

In June 2002, the site has organized a School and Workshop on Models and Algorithms for the Web where Andrei Broder (IBM), Anna Karlin (University of Washington), Bruce Maggs (Akamai Technologies and Carnegie Mellon University), Christos Papadimitriou (California University at

Berkeley) and Prabhakar Raghavan (Verity Inc.) have presented a broad overview of the main research directions in the field.

In September 2002 the site has organized the main European event in algorithmics: the federated conference ALGO 2002 in which ESA02, APPROX02 and WABI02 were colocated. Giorgio Ausiello (as Organizing Committee Chairman, Stefano Leonardi, Fabrizio d'Amore, Paolo Franciosa, Alberto Marchetti and Camil Demetrescu have all actively worked at the organization of the conference.

In 2003 (June 22–27), Camil Demetrescu and Pino Italiano have organized a one week Seminar in Bertinoro, devoted to Algorithms and Data Structures.

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`.

Sixteen ALCOM-FT reports have been authored or co-authored by members of the site in the third project period. Some of these have already 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 Center for Algorithmic Systems (formerly: the Algorithm Design group) of the Institute of Information and Computing Sciences. The algorithms group is involved in WP 3 and 4. The group consists of 7 faculty members/researchers and 4 PhD students and is lead by Professor Jan van Leeuwen. One PhD student (Peter Lennartz) is funded fully by the ALCOM project. He came to us after he had completed his Diplom with ALCOM partner Köln.

In the second project year Dr K.M.J. de Bontridder worked in the group as a postdoc on issues in WP 3.1 (Scheduling and Planning). As planned he left at the beginning of the third project year for a job as algorithm specialist at Siemens VDO Dayton. During the third project year, Peter Bosman finished his Phd thesis on the "Design and application of iterated density-estimation evolutionary algorithms" in the Center. During the third year and the extension the group continued its active cooperations with the National Aerospace Laboratory (NLR) and Philips Research on transportation and optimization problems. For example, Jan van Leeuwen co-supervised the PhD research by Wil Michiels on the analysis of the Karmarkar-Karp differencing method and gave an invited lecture in the 1st Philips Symposium on Intelligent Algorithms (December 11-12, 2002). Excellent contacts are maintained throughout the ALCOM-FT consortium, especially with partners Aarhus, Köln, Patras, Paderborn, and Rome during this period.

During the third year and the extension the research in WP 3 progressed steadily, with excellent progress especially in the work on the no-wait job shop scheduling problem (WP 3.1), multicriteria scheduling (WP 3.1), the analysis of structured networks (WP 3.2), and initializing studies on design and routing problems for mobile ad hoc networks. In the case of no-wait job shop scheduling, a considerable effort was invested in the design and development of a software prototype for suitable instances of the problem. The software prototype took more time than anticipated but was delivered successfully during 2003. In the meantime the prototype is serving as a highly valuable research tool. Much research was also spent on scheduling problems in the case of stochastic processing times (Hoogeveen, Van den Akker). In WP 3.2 the group studied a variety of problems for networks of bounded treewidth such as the 'network reliability problem' (research by Thomas Wolle). Within WP 4 the group has continued to work on new optimization methods, in particular on the use of constraint programming techniques in the (no-wait) job shop scheduling problem. Within WP 4 the group continued its work on computational models for 'evolving' systems and on a suitable complexity theory for these models (research by Peter Verbaan).

Members of the group have authored or co-authored the following ALCOM-FT reports during the third project year and the extension: TR-03-68, TR-03-121, TR-03-122, TR-03-123, TR-03-124, TR-03-188, TR-03-189, TR-03-190, TR-03-191, TR-03-192, TR-03-193, and TR-03-195. The group successfully completed the promised deliverables D19, D27, and D28. The problem instance

database of D28 is for internal use only because of its industrial origin. Papers based on the group's work in scheduling and planning appeared in the Air Traffic Control Quarterly, the INFORMS Journal on Computing, Mathematical Programming B, and the Operations Research Letters. Papers on the group's work in graphs and structured networks appeared e.g. in Algorithmica and Discrete Applied Mathematics. Gerard Tel's book "Introduction to Distributed Algorithms" appeared in a special Foreign Edition in China.

Members of the group have participated actively in the program committees of a number of international conferences and workshops, including COCOON'02, ICALP'02, TCS'02, MFCS'03, SIROCCO'03 and SOFSEM'03. The Utrecht group hosted WG 2003, the 29th Int'l Workshop on Graph-Theoretic Concepts in Computer Science. This very successful workshop was chaired by Hans Bodlaender and attracted over 70 participants. The WG'2003 proceedings edited by Hans Bodlaender recently appeared as LNCS Volume 2880. Han Hoogeveen served on the program committee of the 1st Multidisciplinary International Conference on Scheduling : Theory and Applications (MISTA 2003). Several of the reports and papers of the group were accepted at leading conferences and workshops such as ISAAC'2002 (paper by Thomas Wolle), STACS'2003 (paper by Wil Michiels et al), SWAT'2002, ESA'2002, and WG'02 (papers by Hans Bodlaender et al), all with partial ALCOM-FT support.

Warwick

Warwick contributed to Tasks 1.1, 1.2, 2.4, 2.5, 4.1, 4.3, and 6.2 in the third project period. The research group of the site consists of eight persons: Leslie Goldberg, Paul Goldberg, Steven Kelk, Andrei Krokhin, Russ Martin, Nick Palmer, Mike Paterson, and Alexandre Tiskin. Russ Martin was supported by ALCOM-FT as a Research Fellow for approximately 2 months of this final period. Stephen Kelk passed his PhD viva in November 2003.

Members of the Warwick ALCOM-FT team made research visits to Edinburgh University, the University of New Mexico, AT&T Labs, DIMACS, Yale University, Otago University, Tata Research Development and Design Centre, Leeds University, and Simon Fraser University during the reporting period. We also attended and presented talks at three Dagstuhl Seminars, STOC, the Symposium on Temporal Reasoning, the Symposium on High-level Parallel Programming, the Newton Institute Programme on Combinatorics, Probability and Computing, ISMP, EuroComb, FOCS, STACS, Midland Probability Day, the British Colloquium on Theoretical Computer Science, RAND-APX meetings, ISMVL, Liverpool Algorithms Day, the French Workshop on Complexity and Logic, the British Combinatorial Conference, and an Oberwolfach Seminar. We gave colloquia at Edinburgh University, the University of Bath, Oxford University and Simon Fraser University.

We organised *Warwick Algorithms and Complexity Day* in May 2002 (with speakers: Claire Kenyon, Uri Zwick, Muthu Muthukrishnan, Iain Stewart, Tomasz Radzik, Richard Brent, Leszek Gasieniec, and Vladimir Deineko). We hosted and organised the ALCOM-FT Review Workshop in July 2002. We co-organised the Durham London Mathematical Society Symposium on Markov chains in July 2003.

In this final period, Warwick produced 18 ALCOM-FT Reports of which at least 12 have appeared or are about to appear in conference proceedings or journals.

Cyprus

The activity at site Cyprus is based on the theory group at the Department of Computer Science, University of Cyprus. The group currently comprises three associate/assistant professors, two special scientists, and four students.

During the third year, the site has been very active in the fields of *Algorithmic Game Theory, Distributed Computing*, and *Formal Methods for Specification and Verification of Algorithms*. In

particular, the site has been involved in: Work Package 2 (tasks 2.2, 2.4, 2.5), Work Package 4 (tasks 4.1, 4.3), and Work Package 5 (task 5.2).

In 2002, three Master's students (Ippolytou, Pavlou, and Lazarides) joined the group. Ippolytou and Pavlou have been working on algorithmic and combinatorial properties of *Nash equilibria*; both are expected to graduate in January 2004. Pavlou completed successfully the defense for her Master´s Thesis (entitled *Proof of a Special Case of the Quadratic Fully Mixed Nash Equilibrium Conjecture*) on 01/12 2003. Lazarides has been working on a Master's thesis on *Network Creation Games*. In 2003, Master´s student Eliades joined the group; he has been working on bisimulation problems in stochastic process algebras, and he is expected to graduate around March 2004.

The Cyprus site has maintained strong cooperation links to the CTI and Paderborn sites. This is demonstrated by the numerous visits that were mutually paid and the numerous joint publications. *Algorithmic Game Theory* and *Distributed Computing* have been the main fields on which the cooperation focused. The Cyprus site has also maintained tight cooperation links to the Department of Computer Science, Rensselaer Polytechnic Institute (with Costas Busch and Malik-Magdon Ismail), and to the Department of Computer and Information Sciences, University of Pennsylvania (with Insup Lee).

The following visits were paid by people from other sites of ALCOM-FT to Cyprus: Paul Spirakis (site CTI, 06/08–11/08 2002), Dimitrios Koukopoulos (site CTI, 15/09–15/10 2002), Dimitrios Koukopoulos (site CTI, 13/11–13/12 2002), and Burkhard Monien (site Paderborn, 20/09–05/10 2003). From the Cyprus site, M. Mavronicolas paid the following visits to other sites of ALCOM-FT: Paderborn (24/08–07/09 2002, 21/11–28/11 2002, 07/03–11/03 2003, 22/03–25/03 2003, 05/07–02/08 2003, 05/12–21/12 2003) and CTI (19/12 – 24/12 2002, 27/12 – 30/12 2003)

The following external international collaborators visited site Cyprus where they held short-term research positions: Foto Afrati (National Technical University of Athens, Greece, November 2002) Panagiota Fatourou (University of Ioannina, Greece, March and April 2003) Bernhard Steffen (University of Dortmund, Germany, June 2003) Costas Busch (Rensselaer Polytechnic Institute, USA, August 2003) Malik-Magdon Ismail (Rensselaer Polytechnic Institute, USA, August 2003). From site Cyprus, Anna Philippou visited University of Pennsylvania, USA (June 2003).

M. Mavronicolas gave tutorials and invited talks at the following places: the conference *Euro-Par 2002*, National Technical University of Athens, University of Athens, University of Paderborn (twice), and the conference *ICTCS 2003*. M. Mavronicolas was elected to the program committee of the conferences *SIROCCO 2002*, *ICDCS 2002*, *Euro-Par 2002*, *SPAA 2003*, *IPDPS 2004*, and *SIROCCO 2004*, and to the steering committee of the conference *DISC*. A. Philippou was elected to the program committee of the conferences *Monterey Workshop 2003* and *TACAS 2004*.

Together with Samson Abramsky (Oxford University, UK), M. Mavronicolas is currently co-editing a special volume of the journal *Theoretical Computer Science* on the interaction between Game Theory and Theoretical Computer Science. The title of the volume is *"Game Theory Meets Theoretical Computer Science"*. M. Mavronicolas serves also on the Editorial Board of the journals *Theoretical Computer Science* and *Journal of Interconnection Networks,* and he is the Distributed Computing Column Editor for the *Bulletin of the European Association for Theoretical Computer Science*.

In the third project period, site Cyprus achieved a wide variety of research publications in top journals and conferences of Theoretical Computer Science, such as *Theory of Computing Systems, Theoretical Computer Science, Parallel Processing Letters* and *MFCS, STACS, DISC,* respectively. The research publication record of the site counts (over the reported period) 5 journal papers (with 3 of them in Special Issues of prestigious journals) and 11 conference papers (which either appeared or were accepted). Eleven ALCOM-FT reports have been authored or co-authored by members of the site [57–59, 70, 132, 133, 157–159, 171, 172]. All these have either been published, or have been accepted for publication, or are currently being submitted for publication.

The *Greek Open University* published in 2003 the textbook *"Graph Theory"* (in Greek) by M. Mavronicolas. This is the first Greek textbook on graph theory that is suitable for distant-learning. Moreover, the *Greek Open University* is now about to publish the textbook *"Distributed Systems"* (in Greek) by M. Mavronicolas. This is again the first Greek textbook on distributed

systems that is suitable for distant-learning. In addition, a research monograph entitled *"Counting Networks"* by C. Busch and M. Mavronicolas is expected to appear in the summer of 2004 (*Kluwer* is the publisher); the monograph will be devoted to counting networks, a novel distributed/concurrent data structure to solve problems of counting and coordination in a distributed system. Finally, M. Mavronicolas and P. Spirakis have co-signed (with *Springer*) a contract for a research monograph entitled *"Algorithmic Game Theory"*; the monograph is expected to appear in 2005. The latter monograph is expected to address cooperation, non-cooperation, antagonism and equilibria among dynamic autonomous entities at a suitable proportion and level, and it will provide the first comprehensive coverage of the emerging discipline of algorithmic game theory; the monograph will be part of the series *Texts in Theoretical Computer Science: An EATCS Series* of *Springer*. Finally, A. Philippou together with I. Lee and O. Sokolsky, have been preparing a research monograph with working title *"A Family of Resource-Bounded Process Algebras for Real-Time and Embedded Systems"*.

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 the third project period, as well as the accumulated figures for the entire project. Only non-zero entries are shown.

Third Period

Task	Aarhus		Barcelona		Cologne		INRIA		MPI		Paderborn		CTI		Rome		Utrecht		Warwick		Cyprus		*Total*	
	E	A	E	A	E	A	E	A	E	A	E	A	E	A	E	A	E	A	E	A	E	A	E	A
1.1									4	8	4	3			1	2			2	1.5			11	14.5
1.2	3	4					0.5	1	9	17			4	4					1	1			17.5	27
1.3	1	1	3	2			0.5	2															4.5	5
1.4			3	3			0.5	2							2	5							5.5	10
1.5									4	8													4	8
2.1													3	3									3	3
2.2									3	4			3	3									6	7
2.3									2	2			2	1									4	3
2.4				2			1	3	1	2			2	1					3	3	3	3	10	14
2.5								1			4	2			5	5			3	4	3	3	15	15
2.6									1	1											4	5	5	6
2.7											3	1.5											3	1.5
3.1				1							2	1					11	24					13	26
3.2					3	3					2	1					1.5	4					6.5	8
3.3											3	3					0.5	1.5					3.5	4.5
4.1	4	4	4	7			1.5	18	6	9			3	3					6.5	6	2	2	27	49
4.2	3	4			13	12			4	7	1.5				1	6	3	3					25.5	32
4.3	4	4	3	3			0.5	3			2	1	3	2	1	4		3	2	2	3	2	18.5	24
5.1	4	4	3						3	5	4	5	2	2									16	16
5.2			2					1	3	4			2	2									7	7
5.3					4	3	1	3															5	6
5.4															2	3							2	3
5.5									1	2					2	3							3	5
6.1			2						3	2	1	1	1	1	0.5	0.5	1	1.5					8.5	6
6.2	3	4	2.5	2	1	1	0.5	1	2	1			1	2	0.5	0.5	1	0.5	1	1	1	1	13.5	14
Sum	22	25	22.5	20	21	19	6	35	46	72	26.5	18.5	26	24	15	29	18	37.5	18.5	18.5	16	16	237.5	314.5

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.

Entire Project

Task	Aarhus		Barcelona		Cologne		INRIA		MPI		Paderborn		CTI		Rome		Utrecht		Warwick		Cyprus		*Total*	
	E	A	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	10	15	10	10			3	4			3	3.5			26.5	33.5
1.2	9	11					1.5	4	14	28			10	10					6	6			40.5	59
1.3	1	1	6	6			1	3		6					4	2			1	0.5			13	18.5
1.4			9	9			1	3		0.5					2	5							12	17.5
1.5									6	12													6	12
2.1													9	9	1	1							10	10
2.2									3	7			11	12	2.5	5							16.5	24
2.3									5	7			4	4									9	11
2.4		2		3			3	10	5	8			6	5					12	9	3	3	29	40
2.5							1	4	3	3	14	14			7.5	10			8	11	5	5	38.5	47
2.6									4	4	4	3.5			1	2					4	5	13	14.5
2.7											9	9.5											9	9.5
3.1	2	3		1							4	3					38	45					44	52
3.2					10	10					3	4					6.5	8					19.5	22
3.3											13	13					2.5	3					15.5	16
4.1	10	12	14	18			7.5	50	15	29			9	10					18.5	23	3	3	77	145
4.2	10	12			37	36			9	14	3.5	4	2	2	5	18	8	7					74.5	93
4.3	10	12	9	10			2.5	7	6	12	6.5	6	9	8	4	13	2	6	8.5	9	3	2	60.5	85
5.1	10	10	3						8	10.5	10	10	6	6	1	1							38	37.5
5.2			2				1	3	5	4.5			8	8									16	15.5
5.3			10	10	10	9	3	8	1	4													24	31
5.4				0.5					1	1					6	12							7	13.5
5.5									5	6					5	8							10	14
6.1	1	1	2						6	5	1.5	2.5	3	2	1.5	1.5	1	2.5					16	14.5
6.2	9	12	5	3.5	3	3	2	3	4	3	1.5	1.5	3	4	1.5	1.5	2	1.5	3	3	1	1	35	37
Sum	62	76	60	61	60	58	24	96	110	179.5	80	81	80	80	45	84	60	73	60	65	19	19	660	872.5

Part IV Scientific Reports

In the following, we list all 206 scientific reports (and four further ALCOM-FT related references [2, 136,179,200]) which have been published in the ALCOM-FT Technical Report Series in the period June 1, 2002 to November 31, 2003. All reports are available online at `www.brics.dk/ALCOM-FT/TR`.

References

[1] M. Adler, P. Berenbrink, T. Friedetzky, L. A. Goldberg, P. Goldberg, and M. Paterson. A Proportionate Fair Scheduling Rule with Good Worst-case Performance. Technical Report ALCOMFT-TR-03-48, ALCOM-FT, Warwick, September 2003. Appearance: In *Proceedings of ACM SPAA*, 2003.

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

[3] Marjan van den Akker and Han Hoogeveen. Minimizing the number of tardy jobs. Technical Report ALCOMFT-TR-03-191, ALCOM-FT, Utrecht, December 2003. Appearance: To appear as Chapter 12 in J.Y.-T. Leung (ed.) Handbook of Scheduling: Algorithms, Models, and Performance Analysis.

[4] Ernst Althaus, Fritz Eisenbrand, Stefan Funke, and Kurt Mehlhorn. Point Containment in the Integer Hull of a Polyhedron. Technical Report ALCOMFT-TR-03-184, ALCOM-FT, MPI, December 2003. Appearance: In *Proc. 15th Symposium on Discrete Algorithms (SODA-04)*, New Orleans, USA, 2004. to appear.

[5] Maria Andreou, Sotiris Nikoletseas, and Paul Spirakis. Algorithms and Experiments on Colouring Squares of Planar Graphs. Technical Report ALCOMFT-TR-03-169, ALCOM-FT, Patras, December 2003. Appearance: In *Proceedings of the 2nd Workshop on Efficient and Experimental Algorithms – WEA 2003*, 2003.

[6] Maria Andreou and Paul Spirakis. Efficient Colouring of Squares of Planar Graphs. Technical Report ALCOMFT-TR-03-171, ALCOM-FT, Patras, December 2003.

[7] Thanasis Antoniou, Azzedine Boukerche, Ioannis Chatzigiannakis, George Mylonas, and Sotiris Nikoletseas. A New Energy Efficient and Fault-tolerant Protocol for Data Propagation in Smart Dust Networks using Varying Transmission Range. Technical Report ALCOMFT-TR-03-176, ALCOM-FT, Patras, December 2003. Appearance: In *Proceedings of the 37th ACM/IEEE Annual Simulation Symposium – ANSS 2004*, 2003.

[8] Albert Atserias and María Luisa Bonet. On the Automatizability of Resolution and Related Propositional Proof Systems. Technical Report ALCOMFT-TR-03-104, ALCOM-FT, Barcelona, November 2003. Appearance: *Information and Computation*, 2003. To appear.

[9] Albert Atserias and Víctor Dalmau. A Combinatorial Characterization of Resolution Width. Technical Report ALCOMFT-TR-03-81, ALCOM-FT, Barcelona, November 2003. Appearance: In *18th IEEE Conference on Computational Complexity*, pages 239–246, 2003.

[10] Giorgio Ausiello, Cristina Bazgan, Marc Demange, and Vangelis Th. Paschos. Completeness in differential approximation classes. Technical Report ALCOMFT-TR-03-38, ALCOM-FT, Rome, September 2003. Appearance: In *Mathematical Foundations of Computer Science (MFCS) 2003.* to appear.

[11] Giorgio Ausiello, Paolo G. Franciosa, and Giuseppe F. Italiano. Fully dynamic maintenance of t-spanners on general graphs. Technical Report ALCOMFT-TR-03-23, ALCOM-FT, Rome, September 2003.

[12] Jaume Baixeries and Jose Balcazar. Discrete Deterministic Data Mining as Knowledge Compilation. Technical Report ALCOMFT-TR-03-201, ALCOM-FT, Barcelona, December 2003. Appearance: In *Workshop Discr Math and Data Mining at SIAM DM Conference*, 2003.

[13] Jaume Baixeries, Gemma Casas-Garriga, and Jose Luis Balcazar. A Best-First Strategy For Finding Frequent Sets. Technical Report ALCOMFT-TR-03-198, ALCOM-FT, Barcelona, December 2003. Appearance: In *1st. Journées francophones d'Extraction et de Gestion des Connaissances*, pages 101–107, 2002.

[14] Viviane Baladi and Brigitte Vallée. Euclidean Algorithms are Gaussian. Technical Report ALCOMFT-TR-03-14, ALCOM-FT, INRIA, July 2003.

[15] Lali Barrière, Pierre Fraigniaud, Nicola Santoro, and Dimitrios M. Thilikos. Connected and Internal Graph Searching. Technical Report ALCOMFT-TR-03-69, ALCOM-FT, Barcelona, October 2003. Appearance: In *29th Workshop on Graph Theoretic Concepts in Computer Science (WG 2003)*. Springer, 2003.

[16] Wilhelm Barth, Petra Mutzel, and Michael Jünger. Simple and Efficient Bilayer Crosscounting. Technical Report ALCOMFT-TR-03-94, ALCOM-FT, Cologne, November 2003.

[17] Holger Bast, Kurt Mehlhorn, Guido Schäfer, and Hisao Tamaki. A Heuristic for Dijkstra's Algorithm with Many Targets and its Use inWeighted Matching Algorithms. Technical Report ALCOMFT-TR-03-114, ALCOM-FT, MPI, November 2003. Appearance: *Algorithmica*, 36(1):75–88, 2003.

[18] Holger Bast, Kurt Mehlhorn, Guido Schäfer, and Hisao Tamaki. Matching Algorithms are Fast in Sparse Random Graphs. Technical Report ALCOMFT-TR-03-111, ALCOM-FT, MPI, November 2003. Appearance: In *Conference Proceedings of the 21st International Symposium on Theoretical Aspects of Computer Science (STACS)*, 2004. To appear.

[19] L. Becchetti, J. Könemann, and S. Leonardi. Sharing the cost more efficiently: Improved Approximation for Multicommodity Rent-or-Buy. Technical Report ALCOMFT-TR-03-103, ALCOM-FT, Rome, November 2003.

[20] Luca Becchetti. Modeling Temporal Locality in Paging: A tighter Analysis of LRU and FWF. Technical Report ALCOMFT-TR-03-72, ALCOM-FT, Rome, November 2003.

[21] Luca Becchetti, Stefano Leonardi, Alberto Marchetti-Spaccamela, and Kirk Pruhs. Semiclairvoyant Scheduling. Technical Report ALCOMFT-TR-03-22, ALCOM-FT, Rome, September 2003. Appearance: In *Proc. of the 11th annual European Symposium on Algorithms*, 2003. to appear.

[22] Luca Becchetti, Stefano Leonardi, Alberto Marchetti-Spaccamela, Guido Schäfer, and Tjark Vredeveld. Average Case and Smoothed Competitive Analysis of the Multi-Level Feedback Algorithm. Technical Report ALCOMFT-TR-03-36, ALCOM-FT, Rome and MPI, September 2003. Appearance: In *Proc. of the 44th annual IEEE Symposium on Foundations of Computer Science*, pages 462 – 471, 2003.

[23] Luca Becchetti, Stefano Leonardi, Alberto Marchetti-Spaccamela, Guido Schäfer, and Tjark Vredeveld. Scheduling to minimize flow time metrics. Technical Report ALCOMFT-TR-03-109, ALCOM-FT, Rome and MPI, November 2003. Appearance: In *3rd Workshop on Wireless, Mobile and Ad Hoc Networks (WMAN), satellite workshop of 17th IEEE International Parallel and Distributed Processing Symposium (IPDPS)*.

[24] Rene Beier, Artur Czumaj, Piotr Krysta, and Berthold Vöcking. Computing Equilibria for Congestion Games

with (Im)perfect Information. Technical Report ALCOMFT-TR-03-59, ALCOM-FT, MPI, October 2003. Appearance: In *Proc. 15th Symposium on Discrete Algorithms (SODA-04)*, New Orleans, USA, 2004. to appear.

[25] Rene Beier and Berthold Vöcking. Probabilistic Analysis of Knapsack Core Algorithms. Technical Report ALCOMFT-TR-03-60, ALCOM-FT, MPI, October 2003. Appearance: In *Proc. 15th Symposium on Discrete Algorithms (SODA-04)*, New Orleans, USA, 2004. to appear.

[26] Rene Beier and Berthold Vöcking. Random Knapsack in Expected Polynomial Time. Technical Report ALCOMFT-TR-03-62, ALCOM-FT, MPI, October 2003. Appearance: In *Proc. 35rd Annual ACM Symposium on Theory of Computing (STOC-03)*, pages 232–241, San Diego, USA, 2003. Association of Computing Machinery (ACM), ACM.

[27] Nadia Ben Azzouna, Christine Fricker, and Fabrice Guillemin. Modeling ADSL traffic on an IP backbone link. Technical Report ALCOMFT-TR-03-20, ALCOM-FT, INRIA, August 2003.

[28] Michael A. Bender, Gerth Stølting Brodal, Rolf Fagerberg, Dongdong Ge, Simai He, Haodong Hu, John Iacono, and Alejandro López-Ortiz. The Cost of Cache-Oblivious Searching. Technical Report ALCOMFT-TR-03-76, ALCOM-FT, Aarhus, November 2003. Appearance: In *Proc. 44th Annual Symposium on Foundations of Computer Science*, pages 271–282, 2003.

[29] P. Berenbrink, L. A. Goldberg, P. Goldberg, and R. Martin. Utilitarian Resource Assignment. Technical Report ALCOMFT-TR-03-86, ALCOM-FT, Warwick, November 2003.

[30] Stefan Bertels and Torsten Fahle. A Hybrid Setup for a Hybrid Scenario: Combining Heuristics for the Home Health Care Problem. Technical Report ALCOMFT-TR-03-149, ALCOM-FT, Paderborn, December 2003. Appearance: In *Proceedings of the CP-AI-OR'03*, 2003.

[31] Marcin Bienkowski, Miroslaw Korzeniowski, and Harald Räcke. A Practical Algorithm for Constructing Oblivious Routing Schemes. Technical Report ALCOMFT-TR-03-116, ALCOM-FT, Paderborn, December 2003. Appearance: In *Proceedings of the fifteenth annual ACM symposium on Parallel algorithms and architectures*, pages 24–33. ACM Press, 2003.

[32] M. Blesa and C. Blum. Ant Colony Optimization for the maximum disjoint paths problem. Technical Report ALCOMFT-TR-03-105, ALCOM-FT, Barcelona, November 2003. Appearance: Submitted.

[33] C. Blum and M. Blesa. New Metaheuristic Approaches for the Edge-Weighted k-Cardinality Tree Problem. Technical Report ALCOMFT-TR-03-37, ALCOM-FT, Barcelona, September 2003. Appearance: *Computers & Operations Research.* Accepted for publication, to appear in 2004.

[34] Hans L. Bodlaender, Hajo J. Broersma, Fedor V. Fomin, Artem V. Pyatkin, and Gerhard J. Woeginer. Radio labeling with pre-assigned frequencies. Technical Report ALCOMFT-TR-03-122, ALCOM-FT, Utrecht, December 2003. Appearance: In Rolf Möhring and Rajeev Raman, editors, *Proceedings of the 10th Annual European Symposium on Algorithms, ESA'2002, Lecture Notes in Computer Science 2461*, pages 211–222. Springer-Verlag, 2002.

[35] Hans L. Bodlaender, Michael R. Fellows, and Dimitrios M. Thilikos. Derivation of algorithms for cutwidth and related graph layout parameters. Technical Report ALCOMFT-TR-03-68, ALCOM-FT, Barcelona and Utrecht, October 2003. Appearance: In *28th International*

Symposium on Mathematical Foundations of Computer Science, August 28 - September 1, 2000, Bratislava, Slovak Republic, pages 239–248. Springer, 2003.

[36] Hans L. Bodlaender and Arie M. C. A. Koster. Safe separators for treewidth. Technical Report ALCOMFT-TR-03-124, ALCOM-FT, Utrecht, December 2003. APPEARANCE: To appear in Proceedings ALENEX'04.

[37] Hans L. Bodlaender, Arie M.C.A. Koster, and Frank van den Eijkhof. Pre-Processing Rules for Triangulation of Probabilistic Networks. Technical Report ALCOMFT-TR-03-121, ALCOM-FT, Utrecht, December 2003. APPEARANCE: In J. Breese and D. Koller, editors, *Proceedings of the 17th Conference on Uncertainty in Artificial Intelligence*, pages 32–39, San Francisco, 2001. Morgan Kaufmann.

[38] Hans L. Bodlaender and Gerard Tel. Rectilinear Graphs and Angular Resolution. Technical Report ALCOMFT-TR-03-188, ALCOM-FT, Utrecht, December 2003.

[39] F. Boerner, A. Bulatov, P. Jeavons, and A. Krokhin. Quantified constraints: Algorithms and complexity. Technical Report ALCOMFT-TR-03-54, ALCOM-FT, Warwick, October 2003. APPEARANCE: In *Proceedings of CSL'03*, volume 2803 of *LNCS*, pages 58–70, 2003.

[40] Vincenzo Bonifaci, Benedetto A. Colombo, Camil Demetrescu, Irene Finocchi, and Luigi Laura. A System for Building Animated Presentations over the Web. Technical Report ALCOMFT-TR-03-49, ALCOM-FT, Rome, September 2003. APPEARANCE: In *Proceedings of the AICCSA 2003 Workshop on Practice and Experience with Java Programming in Education, Tunis*, 2003.

[41] Azzedine Boukerche and Sotiris Nikoletseas. Algorithmic Design for Communication in Mobile Ad hoc Networks. Technical Report ALCOMFT-TR-03-178, ALCOM-FT, Patras, December 2003. APPEARANCE: In *Book Chapter, to appear by Springer Verlag*, 2003.

[42] Azzedine Boukerche and Sotiris Nikoletseas. Protocols for Data Propagation in Wireless Sensor Networks: A Survey. Technical Report ALCOMFT-TR-03-172, ALCOM-FT, Patras, December 2003. APPEARANCE: In *Book Chapter, to appear by Kluwer Academic Publishers*, 2003.

[43] Joan Boyar and Lene M. Favrholdt. The Relative Worst Order Ratio for On-Line Algorithms. Technical Report ALCOMFT-TR-03-10, ALCOM-FT, Aarhus, July 2003. APPEARANCE: In *5th Italian Conference on Algorithms and Complexity*, volume 2653 of *LNCS*, pages 58–69, 2003.

[44] Joan Boyar and Lene M. Favrholdt. The Relative Worst Order Ratio for On-Line Bin Packing Algorithms. Technical Report ALCOMFT-TR-03-21, ALCOM-FT, Aarhus, September 2003.

[45] Joan Boyar, Lene M. Favrholdt, and Kim S. Larsen. The Relative Worst Order Ratio Applied to Paging. Technical Report ALCOMFT-TR-03-32, ALCOM-FT, Aarhus, September 2003.

[46] Jacqueline Boyer, Fabrice Guillemin, Philippe Robert, and Bert Zwart. Heavy tailed **M/G/1**-PS queues with impatience and admission control in packet networks. Technical Report ALCOMFT-TR-03-13, ALCOM-FT, INRIA, July 2003. APPEARANCE: In *Infocomm'2003*, San Francisco, USA, 2003.

[47] Andreas Brandstädt, Hans L. Bodlaender, Dieter Kratsch, Michaël Rao, and Jeremy Spinrad. On Algorithms for (P_5,Gem)-Free Graphs. Technical Report ALCOMFT-TR-03-189, ALCOM-FT, Utrecht, December 2003. APPEARANCE: To appear in proceedings FCT 2003.

[48] André Brinkmann, Friedhelm Meyer auf der Heide, Kay Salzwedel, Christian Scheideler, Mario Vodisek, and Ulrich Rückert. Storage Management as Means to cope with Exponential Information Growth. Technical Report ALCOMFT-TR-03-117, ALCOM-FT, Paderborn, December 2003. APPEARANCE: In *Proceedings of SSGRR 2003*, July 2003.

[49] Gerth Stølting Brodal and Rolf Fagerberg. Lower Bounds for External Memory Dictionaries. Technical Report ALCOMFT-TR-03-75, ALCOM-FT, Aarhus, November 2003. APPEARANCE: In *Proc. 14th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 546–554, 2003.

[50] Gerth Stølting Brodal and Rolf Fagerberg. On the Limits of Cache-Obliviousness. Technical Report ALCOMFT-TR-03-74, ALCOM-FT, Aarhus, November 2003. APPEARANCE: In *Proc. 35th Annual ACM Symposium on Theory of Computing*, pages 307–315, 2003.

[51] Gerth Stølting Brodal, Rolf Fagerberg, Thomas Mailund, Christian N. S. Pedersen, and Derek Phillips. Speeding Up Neighbour-Joining Tree Construction. Technical Report ALCOMFT-TR-03-102, ALCOM-FT, Aarhus, November 2003.

[52] Gerth Stølting Brodal, Rolf Fagerberg, Anna Östlin, Christian N. S. Pedersen, and S. Srinivasa Rao. Computing Refined Buneman Trees in Cubic Time. Technical Report ALCOMFT-TR-03-73, ALCOM-FT, Aarhus, November 2003. APPEARANCE: In *Proc. 3rd Workshop on Algorithms in BioInformatics*, volume 2812 of *Lecture Notes in Computer Science*, pages 259–270. Springer Verlag, Berlin, 2003.

[53] Gerth Stølting Brodal, Rolf Fagerberg, and Kristoffer Vinther. Engineering a Cache-Oblivious Sorting Algorithm. Technical Report ALCOMFT-TR-03-101, ALCOM-FT, Aarhus, November 2003. APPEARANCE: In *Proc. Algorithm Engineering & Experiments (ALENEX)*, 2004. To appear.

[54] Gerth Stølting Brodal, Erik D. Demaine, and J. Ian Munro. Fast Allocation and Deallocation with an Improved Buddy System. Technical Report ALCOMFT-TR-03-3, ALCOM-FT, Aarhus, May 2003.

[55] Christoph Buchheim and Michael Jünger. An Integer Programming Approach to Fuzzy Symmetry Detection. Technical Report ALCOMFT-TR-03-96, ALCOM-FT, Cologne, November 2003.

[56] Stefan Burkhardt and Juha Kärkkäinen. Fast Lightweight Suffix Array Construction and Checking. Technical Report ALCOMFT-TR-03-24, ALCOM-FT, MPI, September 2003. APPEARANCE: In *Proceedings of the 14th Annual Symposium on Combinatorial Pattern Matching*, volume 2676 of *Lecture Notes in Computer Science*, pages 55–69. Springer, 2003.

[57] Costas Busch, Malik-Magdon Ismail, Marios Mavronicolas, and Paul Spirakis. An Analysis of Direct Routing - Why do we need Buffers? Technical Report ALCOMFT-TR-03-133, ALCOM-FT, Patras and Cyprus, December 2003.

[58] Costas Busch, Malik-Mgdon Ismail, Marios Mavronicolas, and Roger Wattenhofer. Greedy $\tilde{O}(C+D)$ Hot-Potato Routing on Trees. Technical Report ALCOMFT-TR-03-139, ALCOM-FT, Cyprus, December 2003.

[59] Costas Busch, Marios Mavronicolas, and Paul Spirakis. The Cost of Concurrent, Low-contention Read-Modify-Write. Technical Report ALCOMFT-TR-03-134, ALCOM-FT, Patras and Cyprus, December 2003.

[60] Jesper Makholm Byskov. Chromatic Number in Time $O(2.4023^n)$ Using Maximal Independent Sets. Technical Report ALCOMFT-TR-03-78, ALCOM-FT, Aarhus, November 2003.

[61] Jesper Makholm Byskov, Bolette Ammitzbøll Madsen, and Bjarke Skjernaa. New Algorithms for Exact Satisfiability. Technical Report ALCOMFT-TR-03-77, ALCOM-FT, Aarhus, November 2003.

[62] Ioannis Caragiannis and Christos Kaklamanis. Approximate Path Coloring with Applications to Wavelength Assignment in WDM Optical Networks. Technical Report ALCOMFT-TR-03-187, ALCOM-FT, Patras, December 2003. Appearance: In *Proceedings of the 21st International Symposium on Theoretical Aspects of Computer Science (STACS '04)*, 2004. To appear.

[63] Ioannis Caragiannis, Christos Kaklamanis, and Panagiotis Kanellopoulos. New Results for Energy–Efficient Broadcasting in Wireless Networks. Technical Report ALCOMFT-TR-02-144, ALCOM-FT, Patras, December 2002. Appearance: In *Proceedings of the 13th International Symposium on Algorithms and Computation – ISAAC 2002*, pages 332–343, 2002.

[64] Ioannis Caragiannis, Christos Kaklamanis, and Panagiotis Kanellopoulos. Energy-Efficient Wireless Network Design. Technical Report ALCOMFT-TR-03-157, ALCOM-FT, Patras, December 2003. Appearance: In *Proceedings of the 14th International Symposium on Algorithms and Computation – ISAAC 2003*, 2003. To appear.

[65] Ioannis Caragiannis, Christos Kaklamanis, and Evi Papaioannou. Simple On-line Algorithms for Call Control in Cellular Networks. Technical Report ALCOMFT-TR-03-168, ALCOM-FT, Patras, December 2003. Appearance: In *Proceedings of the 1st Workshop on Approximation and On-line Algorithms – WAOA 2003*, 2003.

[66] Ioannis Caragiannis, Christos Kaklamanis, Pino Persiano, and Anastasios Sidiropoulos. Fractional and Integral Coloring of Locally-Symmetric Sets of Paths on Binary Trees. Technical Report ALCOMFT-TR-03-167, ALCOM-FT, Patras, December 2003. Appearance: In *Proceedings of the 1st Workshop on approximation and On-line Algorithms – WAOA 2003*, 2003.

[67] Gemma Casas-Garriga. Discovering Unbounded Episodes in Sequential Data. Technical Report ALCOMFT-TR-03-197, ALCOM-FT, Barcelona, December 2003. Appearance: In *7th. European conference on Principles and Practice of knowledge Discovery in Databases*, pages 83–94, 2003.

[68] Gemma Casas-Garriga. Statistical Strategies for Pruning All the Uninteresting Association Rules. Technical Report ALCOMFT-TR-03-199, ALCOM-FT, Barcelona, December 2003.

[69] Gemma Casas-Garriga. Towards a formal Framework for mining Mining General Patterns from Structured Data. Technical Report ALCOMFT-TR-03-196, ALCOM-FT, Barcelona, December 2003. Appearance: In *2nd KDD Int'l Workshop on Multirelational Datamining*, pages 14–26, 2003.

[70] Ioannis Chatzigiannakis, Tassos Dimitriou, Marios Mavronicolas, Sotiris Nikoletseas, and Paul Spirakis. A Comparative Study of Protocols for Efficient Data Propagation in Smart Dust Networks. Technical Report ALCOMFT-TR-03-165, ALCOM-FT, Patras and Cyprus, December 2003. Appearance: In *Proceedings of 9th International International Conference on Parallel and Distributed Computing – EUROPAR 2003. Also accepted in the Parallel Processing (PPL) Journal, to appear in 2003*, 2003. Distinguished Paper.

[71] Ioannis Chatzigiannakis, Tassos Dimitriou, Sotiris Nikoletseas, and Paul Spirakis. A Probabilistic Algorithm for Efficient and Robust Data Propagation in Smart Dust Networks. Technical Report ALCOMFT-TR-03-166, ALCOM-FT, Patras, December 2003. Appearance: In *Proceedings of 5th European Wireless (EW) Conference, to appear in 2004*, 2003.

[72] Ioannis Chatzigiannakis, Elena Kaltsa, and Sotiris Nikoletseas. On the Effect of User Mobility and Density on the Performance of Protocols for Ad-hoc Mobile Networks. Technical Report ALCOMFT-TR-03-164, ALCOM-FT, Patras, December 2003.

[73] Ioannis Chatzigiannakis, Athanasios Kinalis, Athanasios Poulakidas, Grigorios Prasinos, and Christos Zaroliagis. DAP: A Generic Platform for the Simulation of Distributed Algorithms. Technical Report ALCOMFT-TR-03-155, ALCOM-FT, Patras, December 2003.

[74] Ioannis Chatzigiannakis, Michael Markou, and Sotiris Nikoletseas. Distributed Circle Formation for Anonymous Oblivious Robots. Technical Report ALCOMFT-TR-03-163, ALCOM-FT, Patras, December 2003.

[75] Ioannis Chatzigiannakis and Sotiris Nikoletseas. A Sleep-Awake Protocol for Information Propagation in Smart Dust Networks. Technical Report ALCOMFT-TR-03-161, ALCOM-FT, Patras, December 2003. Appearance: In *Proceedings of 17th International Parallel & Distributed Processing Symposium, 3rd International Workshop on Mobile and Ad-hoc Networks (WMAN) – IPDPS 2003 Workshops*. IEEE/ACM, 2003.

[76] Ioannis Chatzigiannakis, Sotiris Nikoletseas, and Paul Spirakis. Distributed Communication Algorithms for Ad-hoc Mobile Networks. Technical Report ALCOMFT-TR-03-179, ALCOM-FT, Patras, December 2003. Appearance: In *Book Chapter in the "Current Trends in Theoretical Computer Science", to appear in 2004*, 2003.

[77] Ioannis Chatzigiannakis, Sotiris Nikoletseas, and Paul Spirakis. Efficient and Robust Protocols for Local Detection and Propagation in Smart Dust Networks. Technical Report ALCOMFT-TR-03-175, ALCOM-FT, Patras, December 2003. Appearance: In *In the ACM Journal of Mobile Networks (MONET) - Special Issue on Parallel Processing Issues in Mobile Computing, to appear in 2003*, 2003.

[78] Frédéric Chyzak. Algorithms Seminar, 2001–2002. Technical Report ALCOMFT-TR-03-87, ALCOM-FT, INRIA, November 2003. Appearance: INRIA Research Report # 5003, November 2003., 2003.

[79] D. Cohen, M. Cooper, P. Jeavons, and A. Krokhin. Identifying efficiently solvable cases of non-Boolean Max CSP. Technical Report ALCOMFT-TR-03-56, ALCOM-FT, Warwick, October 2003. Appearance: In *Proceedings of 21th Symposium on Theoretical Aspects of Computer Science, STACS'04*, 2004. accepted.

[80] D. Cohen, M. Cooper, P. Jeavons, and A. Krokhin. Multimorphisms: Classifying the complexity of soft constraints. Technical Report ALCOMFT-TR-03-55, ALCOM-FT, Warwick, October 2003. Appearance: In *Proceedings of CP'03*, volume 2833 of *LNCS*, 2003. to appear.

[81] Artur Czumaj, Funda Ergun, Lance Fortnow, Avner Magen, Ilan Newman, Ronitt Rubinfeld, and Christian Sohler. Sublinear-Time Approximation of Euclidean Minimum Spanning Tree. Technical Report ALCOMFT-TR-03-92, ALCOM-FT, Paderborn, November 2003. Appearance: In *In Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'03)*, pages 813 – 822, 2003.

[82] Artur Czumaj and Christian Sohler. Abstract Combinatorial Programs and Efficient Property Testers. Technical Report ALCOMFT-TR-03-91, ALCOM-FT, Paderborn, November 2003. Appearance: In *Proceedings of the 43rd Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 83 – 92, 2002.

[83] J. D´ ıaz, M. Serna, and N. Wormald. Computation of the bisection width for random d-regular graphs. Technical Report ALCOMFT-TR-03-44, ALCOM-FT, Barcelona, September 2003. Appearance: In *6th Latin American Theoretical Informatics conference*, Lecture Notes in Computer Science. Springer-Verlag, 2004. to appear.

[84] Valentina Damerow, Harald Räcke, Friedhelm Meyer auf der Heide, Christian Scheideler, and Christian Sohler. Smoothed Motion Complexity. Technical Report ALCOMFT-TR-03-89, ALCOM-FT, Paderborn, November 2003. Appearance: In *Algorithms – Proceedings of the 11th European Symposium on Algorithms (ESA)*, pages 161–171, 2003.

[85] Ivan B. Damgård and Gudmund Skovbjerg Frandsen. An Extended Quadratic Frobenius Primality Test with Average and Worst Case Error Estimates. Technical Report ALCOMFT-TR-03-40, ALCOM-FT, Aarhus, September 2003. Appearance: In *Fundamentals of Computation Theory—FCT 2003 (Malmö, Sweden, August 2003)*, volume 2751 of *Lecture Notes in Comput. Sci.*, pages 118–131. Springer, Berlin, 2003.

[86] Ivan Bjerre Damgård and Gudmund Skovbjerg Frandsen. Efficient Algorithms for gcd and Cubic Residuosity in the Ring of Eisenstein Integers. Technical Report ALCOMFT-TR-03-39, ALCOM-FT, Aarhus, September 2003. Appearance: In *Fundamentals of Computation Theory—FCT 2003 (Malmö, Sweden, August 2003)*, volume 2751 of *Lecture Notes in Comput. Sci.*, pages 109–117. Springer, Berlin, 2003.

[87] Josep Díaz, Mathew Penrose, Jordi Petit, and Maria Serna. Approximating Layout Problems on Random Geometric Graphs. Technical Report ALCOMFT-TR-03-84, ALCOM-FT, Barcelona, November 2003. Appearance: *Journal of Algorithms*, 39(1):78–116, 2001.

[88] Thomas Decker, Thomas Lücking, and Burkhard Monien. A 5/4-approximation algorithm for scheduling identical malleable tasks. Technical Report ALCOMFT-TR-03-148, ALCOM-FT, Paderborn, December 2003. Appearance: In *Proceedings of the 1st Workshop on Approximation and Online Algorithms (WAOA 2003), to appear*, 2003.

[89] Vladimir Deineko and Alexander Tiskin. The Double-Tree Approximation for Metric TSP: Is The Best Good Enough? Technical Report ALCOMFT-TR-03-8, ALCOM-FT, Warwick, June 2003. Appearance: In *Proceedings of ISMP'03*, 2003.

[90] Erik D. Demaine, Fedor V. Fomin, MohammadTaghi Hajiaghayi, and Dimitrios M. Thilikos. Bidimensional Parameters and Local Treewidth. Technical Report ALCOMFT-TR-03-64, ALCOM-FT, Barcelona, October 2003.

[91] Erik D Demaine, MohammadTaghi Hajiaghayi, Fedor V. Fomin, and Dimitrios M. Thilikos. Fixed-Parameter Algorithms for the (k, r)-Center in Planar Graphs and Map Graphs. Technical Report ALCOMFT-TR-03-67, ALCOM-FT, Barcelona, October 2003. Appearance: In *ICALP2003 Thirtieth International Colloquium on Automata, Languages and Programming Eindhoven, The Netherlands*. Springer, 2003.

[92] Roman Dementiev, Lutz Kettner, Jens Mehnert, and Peter Sanders. Engineering a Sorted List Data Structure for 32 Bit Keys. Technical Report ALCOMFT-TR-03-118, ALCOM-FT, MPI, December 2003. Appearance: In *Workshop on Algorithm Engineering & Experiments*, New Orleans, 2004. accepted.

[93] Roman Dementiev and Peter Sanders. Asynchronous Parallel Disk Sorting. Technical Report ALCOMFT-TR-03-30, ALCOM-FT, MPI, September 2003. Appearance: In *15th ACM Symposium on Parallelism in Algorithms and Architectures*, pages 138–148, San Diego, 2003.

[94] Camil Demetrescu, Stefano Emiliozzi, and Giuseppe F. Italiano. Experimental Analysis of Dynamic All Pairs Shortest Path Algorithms. Technical Report ALCOMFT-TR-03-9, ALCOM-FT, Rome, July 2003. Appearance: In *To appear in Proceedings of the 15th ACM-SIAM Symposium on Discrete Algorithms (SODA 2004)*, 2004.

[95] Camil Demetrescu and Irene Finocchi. A Portable Virtual Machine for Program Debugging and Directing. Technical Report ALCOMFT-TR-03-1, ALCOM-FT, Rome, March 2003. Appearance: In *To appear in Proceedings of the 19th ACM Symposium on Applied Computing (SAC 2004), Software Engineering Track. Nicosia, Cyprus, March 14 -17, 2004.*, 2004.

[96] Camil Demetrescu and Irene Finocchi. Combinatorial Algorithms for Feedback Problems in Directed Graphs. Technical Report ALCOMFT-TR-03-19, ALCOM-FT, Rome, August 2003. Appearance: *Information Processing Letters*, 86:129–136, 2003.

[97] J. Díaz, N. Do, M. Serna, and N. Wormald. Bounds on the max and min bisection of random cubic and random 4-regular graphs. Technical Report ALCOMFT-TR-03-43, ALCOM-FT, Barcelona, September 2003. Appearance: *Theoretical Computer Science*, 307:531–547, 2003.

[98] Josep Díaz, Jordi Petit, and Maria Serna. Random scaled sector graphs. Technical Report ALCOMFT-TR-02-140, ALCOM-FT, Barcelona, June 2002.

[99] Josep Diaz, Maria Serna, and Dimitrios M. Thilikos. Counting List H-Colorings and Variants. Technical Report ALCOMFT-TR-03-63, ALCOM-FT, Barcelona, October 2003. Appearance: In *European conference on Combinatorics, Graph Theory and Applications*, 2003.

[100] Tassos Dimitriou, Sotiris Nikoletseas, and Paul Spirakis. The Infection Time of Graphs. Technical Report ALCOMFT-TR-03-183, ALCOM-FT, Patras, December 2003.

[101] Debora Donato, Luigi Laura, Stefano Leonardi, and Stefano Millozzi. Large scale properties of the Webgraph. Technical Report ALCOMFT-TR-03-200, ALCOM-FT, Rome, December 2003. Appearance: *European Journal of Physics B*, 2003. Under review.

[102] Devdatt Dubashi, Luigi Laura, and Alessandro Panconesi. Analysis and Experimental Evaluation of a simple Algorithm for Collaborative Filtering. Technical Report ALCOMFT-TR-03-186, ALCOM-FT, Rome, December 2003. Appearance: In *Proc. of FSTTCS 2003 - THE 23rd CONFERENCE ON FOUNDATIONS OF SOFTWARE TECHNOLOGY AND THEORETICAL COMPUTER SCIENCE, Mumbai (India), December 15–17, 2003.*, 2003.

[103] Philippe Duchon, Philippe Flajolet, Guy Louchard, and Gilles Schaeffer. Boltzmann Sampler for the Random Generation of Combinatorial Structures. Technical Report ALCOMFT-TR-03-15, ALCOM-FT, INRIA, July 2003.

[104] Vincent Dumas, Fabrice Guillemin, and Philippe Robert. Admission control of leaky bucket regulated sources in a queueing system with priority. Technical Report ALCOMFT-TR-03-11, ALCOM-FT, INRIA, July 2003. Appearance: In *ITC'18*, Berlin, September 2003.

[105] Marianne Durand and Philippe Flajolet. Loglog Counting of Large Cardinalities. Technical Report ALCOMFT-TR-03-18, ALCOM-FT, INRIA, July 2003.

[106] Charilaos Efthymiou, Sotiris Nikoletseas, and Jose Rolim. Energy Balanced Data Propagation in Wireless Sensor Networks. Technical Report ALCOMFT-TR-03-177, ALCOM-FT, Patras, December 2003.

[107] Frank van den Eijkhof, Hans L. Bodlaender, and Arie M.C.A. Koster. Safe reduction rules for weighted treewidth. Technical Report ALCOMFT-TR-03-123, ALCOM-FT, Utrecht, December 2003. Appearance: In Luděk Kučera, editor, *Proceedings 28th Int. Workshop on Graph Theoretic Concepts in Computer Science, WG'02*, pages 176–185. Springer-Verlag, Lecture Notes in Computer Science 2344, 2002.

[108] Friedrich Eisenbrand, Stefan Funke, Joachim Reichel, and Elmar Schömer. Packing a Trunk. Technical Report ALCOMFT-TR-03-113, ALCOM-FT, MPI, November 2003. Appearance: In *Algorithms - ESA 2003, 11th Annual European Symposium, Budapest, Hungary, September 2003*, pages 618–629, 2003.

[109] Robert Elsässer, Rastislav Královic, and Burkhard Monien. Sparse Topologies with Small Spectrum Size. Technical Report ALCOMFT-TR-03-150, ALCOM-FT, Paderborn, December 2003. Appearance: *Theoretical Computer Science*, 307:549–565, 2003.

[110] Robert Elsässer, Thomas Lücking, and Burkhard Monien. On Spectral Bounds for the k-Partitioning of Graphs. Technical Report ALCOMFT-TR-03-145, ALCOM-FT, Paderborn, December 2003. Appearance: *Theory of Computing Systems*, 36:461–478, 2003.

[111] Robert Elsässer and Burkhard Monien. Diffusion Load Balancing in Static and Dynamic Networks. Technical Report ALCOMFT-TR-03-151, ALCOM-FT, Paderborn, December 2003. Appearance: In *Proc. Internat. Workshop on Ambient Intelligence Computing*, pages 49–62, 2003.

[112] Robert Elsässer and Burkhard Monien. Load Balancing of Unit Size Tokens and Expansion Properties of Graphs. Technical Report ALCOMFT-TR-03-152, ALCOM-FT, Paderborn, December 2003. Appearance: In *Proc. of the 15th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA 2003)*, pages 266–273, 2003.

[113] Torsten Fahle. Simple and Fast: Improving a Branch-And-Bound Algorithm for Maximum Clique. Technical Report ALCOMFT-TR-03-180, ALCOM-FT, Paderborn, December 2003. Appearance: In *Proc. of the 10th Annual European Symposium on Algorithms (ESA'02)*, LNCS 2461, pages 485–498, 2002.

[114] Torsten Fahle and Meinolf Sellmann. Cost based Filtering for the Constrained Knapsack Problem. Technical Report ALCOMFT-TR-03-182, ALCOM-FT, Paderborn, December 2003. Appearance: *Annals of Operations Research*, 115:73–94, 2003.

[115] Rainer Feldmann, Martin Gairing, Thomas Lücking, Burkhard Monien, and Manuel Rode. Nashification and the Coordination Ratio for a Selfish Routing Game. Technical Report ALCOMFT-TR-03-147, ALCOM-FT, Paderborn, December 2003. Appearance: In *Proceedings of the 30th International Colloquium on Automata, Languages, and Programming (ICALP'03)*, LNCS 2719, pages 514–526, 2003.

[116] Rainer Feldmann, Martin Gairing, Thomas Lücking, Burkhard Monien, and Manuel Rode. Selfish Routing in Non-cooperative Networks: A Survey. Technical Report ALCOMFT-TR-03-146, ALCOM-FT, Paderborn, December 2003. Appearance: In *Proceedings of the 28th International Symposium on Mathematical Foundations of Computer Science (MFCS'03)*, LNCS 2747, pages 21–45, 2003.

[117] James A. Fill, Philippe Flajolet, and Nevin Kapur. Singularity Analysis, Hadamard Products, and Tree Recurrences. Technical Report ALCOMFT-TR-03-17, ALCOM-FT, INRIA, July 2003.

[118] Philippe Flajolet, Joaquim Gabarró, and Helmut Pekari. Analytic Urns. Technical Report ALCOMFT-TR-03-16, ALCOM-FT, Barcelona and INRIA, July 2003.

[119] Philippe Flajolet, Bruno Salvy, and Gilles Schaeffer. Airy Phenomena and Analytic Combinatorics of Connected Graphs. Technical Report ALCOMFT-TR-02-143, ALCOM-FT, INRIA, October 2002.

[120] Philippe Flajolet and Robert Sedgewick. Analytic Combinatorics: Basic Complex Asymptotics. Technical Report ALCOMFT-TR-03-100, ALCOM-FT, INRIA, November 2003.

[121] Philippe Flajolet, Wojciech Szpankowski, and Brigitte Vallée. Hidden Word Statistics. Technical Report ALCOMFT-TR-02-142, ALCOM-FT, INRIA, October 2002.

[122] Fedor V. Fomin and Dimitrios M. Thilikos. A Simple and Fast Approach for Solving Problems on Planar Graphs. Technical Report ALCOMFT-TR-03-106, ALCOM-FT, Barcelona, November 2003. Appearance: In *Proceedings 21st International Symposium on Theoretical Aspects of Computer Science, STACS 2004*. Springer Verlag, Lecture Notes in Computer Science, 2004.

[123] Fedor V. Fomin and Dimitrios M. Thilikos. Dominating sets and local treewidth. Technical Report ALCOMFT-TR-03-66, ALCOM-FT, Barcelona, October 2003. APPEARANCE: In *11th Annual European Symposium on Algorithms. Budapest*, pages 221–229. Springer, 2003.

[124] Fedor V. Fomin and Dimitrios M. Thilikos. Dominating Sets in Planar Graphs: Branch-Width and Exponential Speed-up. Technical Report ALCOMFT-TR-03-71, ALCOM-FT, Barcelona, October 2003. APPEARANCE: In *14th ACM-SIAM Symposium on Discrete Algorithms, SODA 2003*, pages 168–177, 2003.

[125] Fedor V. Fomin and Dimitrios M. Thilikos. New upper bounds on the decomposability of planar graphs and fixed parameter algorithms. Technical Report ALCOMFT-TR-03-65, ALCOM-FT, Barcelona, October 2003.

[126] Dimitris Fotakis. Incremental Algorithms for Facility Location and k-Median. Technical Report ALCOMFT-TR-03-156, ALCOM-FT, Patras, December 2003.

[127] Dimitris Fotakis. On the Competitive Ratio for Online Facility Location. Technical Report ALCOMFT-TR-03-25, ALCOM-FT, MPI, September 2003. APPEARANCE: In *Proceedings of the 30th International Colloquium on Automata, Languages and Programming (ICALP '03)*, volume 2719 of *Lecture Notes in Computer Science*, pages 637–652. Springer, 2003.

[128] Dimitris Fotakis, Rasmus Pagh, Peter Sanders, and Paul Spirakis. Space Efficient Hash Tables With Worst Case Constant Access Time. Technical Report ALCOMFT-TR-03-26, ALCOM-FT, Aarhus and Patras and MPI, September 2003. APPEARANCE: In *Proceedings of the 20th Annual Symposium on Theoretical Aspects of Computer Science (STACS '03)*, volume 2607 of *Lecture Notes in Computer Science*, pages 271–282. Springer, 2003.

[129] Jens S. Frederiksen and Kim S. Larsen. On-Line Seat Reservations via Off-Line Seating Arrangements. Technical Report ALCOMFT-TR-03-31, ALCOM-FT, Aarhus, September 2003. APPEARANCE: In *Eighth International Workshop on Algorithms and Data Structures*, volume 2748 of *Lecture Notes in Computer Science*, pages 174–185. Springer-Verlag, 2003.

[130] Stefan Funke, Domagoj Matijevic, and Peter Sanders. Approximating Energy Efficient Paths in Wireless Multi-Hop Networks. Technical Report ALCOMFT-TR-03-28, ALCOM-FT, MPI, September 2003. APPEARANCE: In *11th European Symposium on Algorithms (ESA)*, 2003. to appear.

[131] J. Gabarró, A. A. Stewart, M. Clint, E. Boyle, and I. Vallejo. Computational models for Web- and Grid-based computation. Technical Report ALCOMFT-TR-03-42, ALCOM-FT, Barcelona, September 2003. APPEARANCE: In *Euro-Par 2003*, pages 640–650. Springer-Verlag, LNCS 2790, 2003.

[132] Martin Gairing, Thomas Lücking, Marios Mavronicolas, and Burkhard Monien. Computing Nash Equilibria for Scheduling on Restricted Parallel Links. Technical Report ALCOMFT-TR-03-159, ALCOM-FT, Paderborn and Cyprus, December 2003. APPEARANCE: In *submitted for publication*, 2003.

[133] Martin Gairing, Thomas Lucking, Marios Mavronicolas, Burkhard Monien, and Paul Spirakis. The Structure and Complexity of Extreme Nash Equilibria. Technical Report ALCOMFT-TR-03-141, ALCOM-FT, Paderborn and Patras and Cyprus, December 2003.

[134] Anna Gal and Peter Bro Miltersen. The cell probe complexity of succinct data structures. Technical Report ALCOMFT-TR-03-80, ALCOM-FT, Aarhus, November 2003. APPEARANCE: In *Proc. ICALP'03*, Lecture Notes in Computer Science, 2003.

[135] Ricard Gavaldà and Denis Thérien. Algebraic Characterizations of Small Classes of Boolean Functions. Technical Report ALCOMFT-TR-03-45, ALCOM-FT, Barcelona, September 2003. APPEARANCE: In *Proc. 20th Intl. Symposium on Theoretical Aspects of Computer Science (STACS'03)*, volume 2607, pages 331–342. Springer-Verlag LNCS, 2003.

[136] 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.

[137] L. A. Goldberg, M. Jerrum, and M. Paterson. The computational complexity of two-state spin systems. Technical Report ALCOMFT-TR-03-46, ALCOM-FT, Warwick, September 2003. Appearance: *Random Structures and Algorithms*, 2003. To appear.

[138] L. A. Goldberg, R. Martin, and M. Paterson. Random sampling of 3-colourings in Z^2. Technical Report ALCOMFT-TR-03-47, ALCOM-FT, Warwick, September 2003. Appearance: *Random Structures and Algorithms*. To appear.

[139] L. A. Goldberg, R. Martin, and M. Paterson. Sampling 10-colourings of the triangular lattice. Technical Report ALCOMFT-TR-03-85, ALCOM-FT, Warwick, November 2003.

[140] Paul Wilfred Goldberg. Bounds for the Convergence Rate of Randomized Local Search in a Multiplayer, Uniform Resource Sharing Game. Technical Report ALCOMFT-TR-03-107, ALCOM-FT, Warwick, November 2003.

[141] Sven Grothklags. Fleet Assignment with Connection Dependent Ground Times. Technical Report ALCOMFT-TR-03-173, ALCOM-FT, Paderborn, December 2003. Appearance: In *Proc. of the 11th European Symposium on Algorithms(ESA'03)*, pages 667–678, 2003.

[142] Matthias Grünewald, Tamás Lukovszki, Christian Schindelhauer, and Klaus Volbert. Distributed Maintenance of Resource Efficient Wireless Network Topologies. Technical Report ALCOMFT-TR-03-127, ALCOM-FT, Paderborn, December 2003. Appearance: In *European Conference on Parallel Computing (EUROPAR'02)*, pages 935–946, 2002.

[143] Fabrice Guillemin, Philippe Robert, and Bert Zwart. Tail Asymptotics for Processor Sharing Queues. Technical Report ALCOMFT-TR-03-12, ALCOM-FT, INRIA, July 2003.

[144] MohammadTaghi Hajiaghayi, Naomi Nishimura, Prabhakar Ragde, and Dimitrios M. Thilikos. Fast approximation schemes for $K_{3,3}$-minor-free or K_5-minor-free graphs. Technical Report ALCOMFT-TR-03-70, ALCOM-FT, Barcelona, October 2003. Appearance: In *Euroconference on Combinatorics, Graph Theory and Applications*, 2001.

[145] Kristoffer Arnsfelt Hansen. Constant width planar computation characterizes ACC^0. Technical Report ALCOMFT-TR-03-194, ALCOM-FT, Aarhus, December 2003. Appearance: In *21st International Symposium on Theoretical Aspects of Computer Science, to appear*, Lecture Notes in Computer Science. Springer, 2004.

[146] Kristoffer Arnsfelt Hansen, Peter Bro Miltersen, and V Vinay. Circuits on Cylinders. Technical Report ALCOMFT-TR-03-79, ALCOM-FT, Aarhus, November 2003. Appearance: In *Proc. FCT'03*, Lecture Notes in Computer Science, 2003.

[147] Han Hoogeveen. Multicriteria scheduling. Technical Report ALCOMFT-TR-03-190, ALCOM-FT, Utrecht, December 2003. Appearance: Invited review: to appear in European Journal of Scheduling.

[148] Csanad Imreh. Scheduling problems on two sets of identical machines. Technical Report ALCOMFT-TR-03-58, ALCOM-FT, MPI, October 2003. Appearance: *Computing*, 70:277–294, 2003.

[149] P. Jonsson and A. Krokhin. Computational complexity of auditing finite attributes in statistical databases. Technical Report ALCOMFT-TR-03-108, ALCOM-FT, Warwick, November 2003.

[150] Michael Jünger, Sebastian Leipert, and Merijam Percan. Triangulating Clustered Graphs. Technical Report ALCOMFT-TR-03-97, ALCOM-FT, Cologne, November 2003.

[151] Alexis Kaporis, Christos Makris, Spyros Sioutas, Athanasios Tsakalidis, Kostas Tsichlas, and Christos Zaroliagis. Improved Bounds for Finger Search on a RAM. Technical Report ALCOMFT-TR-03-154, ALCOM-FT, Patras, December 2003. Appearance: In *Algorithms – ESA 2003*, volume 2832, pages 325–336. Lect. Notes in Comp. Sc., 2003.

[152] Juha Kärkkäinen and Peter Sanders. Simple Linear Work Suffix Array Construction. Technical Report ALCOMFT-TR-03-29, ALCOM-FT, MPI, September 2003. Appearance: In *30th International Colloquium on Automata, Languages and Programming*, number 2719 in LNCS, pages 943–955, 2003.

[153] Irit Katriel and Sven Thiel. Fast Bound Consistency for the Global Cardinality Constraint. Technical Report ALCOMFT-TR-03-27, ALCOM-FT, MPI, September 2003. Appearance: In To appear, editor, *Proceedings of the 9th International Conference on Principles and Practice of Constraint Programming (CP 2003)*, 2003.

[154] Jan Klein and Gabriel Zachmann. ADB-Trees: Controlling the Error of Time-Critical Collision Detection. Technical Report ALCOMFT-TR-03-132, ALCOM-FT, Paderborn, December 2003. Appearance: In *8th International Fall Workshop Vision, Modeling, and Visualization (VMV)*, pages 37–45, University München, Germany, November19–21 2003.

[155] Jan Klein and Gabriel Zachmann. Time-Critical Collision Detection Using an Average-Case Approach. Technical Report ALCOMFT-TR-03-131, ALCOM-FT, Paderborn, December 2003. Appearance: In *Proc. ACM Symp. on Virtual Reality Software and Technology (VRST)*, pages 22–31, Osaka, Japan, October1–3 2003.

[156] Jens S. Kohrt and Kirk Pruhs. A constant approximation algorithm for sorting buffers. Technical Report ALCOMFT-TR-03-160, ALCOM-FT, Aarhus, December 2003.

[157] Dimitrios Koukopoulos, Marios Mavronicolas, Sotiris Nikoletseas, and Paul Spirakis. The Impact of Network Structure on the Stability of Greedy Protocols. Technical Report ALCOMFT-TR-03-137, ALCOM-FT, Patras and Cyprus, December 2003.

[158] Dimitrios Koukopoulos, Marios Mavronicolas, and Paul Spirakis. Instability of Networks with Quasi-Static Link Capacities. Technical Report ALCOMFT-TR-03-135, ALCOM-FT, Patras and Cyprus, December 2003.

[159] Dimitrios Koukopoulos, Marios Mavronicolas, and Paul Spirakis. Performance and Stability Bounds for Dynamic Networks. Technical Report ALCOMFT-TR-03-136, ALCOM-FT, Patras and Cyprus, December 2003.

[160] Annamaria Kovacs. Sum-Multicoloring on Paths. Technical Report ALCOMFT-TR-03-112, ALCOM-FT, MPI, November 2003. Appearance: In *Proc. 21st International Symposium on Theoretical Aspects of Computer Science (STACS 2004)*, LNCS. Springer, 2004, to appear.

[161] A. Krokhin, P. Jeavons, and P. Jonsson. Constraint satisfaction problems on intervals and lengths. Technical Report ALCOMFT-TR-03-51, ALCOM-FT, Warwick, October 2003. Appearance: *SIAM Journal on Discrete Mathematics*. to appear.

[162] A. Krokhin, P. Jeavons, and P. Jonsson. Reasoning about temporal relations: The maximal tractable subalgebras of Allen's interval algebra. Technical Report ALCOMFT-TR-03-52, ALCOM-FT, Warwick, October 2003. Appearance: *Journal of the ACM*, 50(5):591–640, 2003.

[163] A. Krokhin and P. Jonsson. Recognizing frozen variables in constraint satisfaction problems. Technical Report ALCOMFT-TR-03-57, ALCOM-FT, Warwick, October 2003.

[164] A. Krokhin and B. Larose. Solving order constraints in logarithmic space. Technical Report ALCOMFT-TR-03-53, ALCOM-FT, Warwick, October 2003. Appearance: In *Proceedings of STACS'03*, volume 2607 of *LNCS*, pages 379–390, 2003.

[165] Luigi Laura, Stefano Leonardi, Stefano Millozzi, Ulrich Meyer, and Jop Sibeyn. Algorithms and Experiments for the Webgraph. Technical Report ALCOMFT-TR-03-120, ALCOM-FT, Rome and MPI, December 2003. Appearance: In *11th Annual Symposium on Algorithms*, pages 703–714. Springer-Verlag, LNCS 2832, 2003.

[166] Luigi Laura, Umberto Nanni, and Fabiano Sarracco. Query Transformation through approximated LSI computation. Technical Report ALCOMFT-TR-03-185, ALCOM-FT, Rome, December 2003.

[167] Stefano Leonardi and Guido Schäfer. Cross-Monotonic Cost-Sharing Methods for Connected Facility Location Games. Technical Report ALCOMFT-TR-03-119, ALCOM-FT, Rome and MPI, December 2003.

[168] Magali Lescot and Mireille Régnier. Motif statistics on plants datasets. Technical Report ALCOMFT-TR-03-99, ALCOM-FT, INRIA, November 2003.

[169] Eythan Levy, Guy Louchard, and Jordi Petit. A distributed algorithm to find Hamiltonian cycles in $G_{n,p}$ random graphs. Technical Report ALCOMFT-TR-03-82, ALCOM-FT, Barcelona, November 2003.

[170] Frauke Liers, Michael Jünger, Gerhard Reinelt, and Giovanni Rinaldi. Computing Exact Ground-States of Hard Ising Spin-Glass. Technical Report ALCOMFT-TR-03-95, ALCOM-FT, Cologne, November 2003.

[171] Thomas Lucking, Marios Mavronicolas, Burkhard Monien, and Manuel Rode. A New model for Selfish Routing. Technical Report ALCOMFT-TR-03-138, ALCOM-FT, Paderborn and Cyprus, December 2003.

[172] Thomas Lucking, Marios Mavronicolas, Burkhard Monien, Manuel Rode, Paul Spirakis, and Imrich Vrto. Which is the Worst-case Nash Equilibrium? Technical Report ALCOMFT-TR-03-140, ALCOM-FT, Paderborn and Patras and Cyprus, December 2003.

[173] Conrado Martínez. Partial Quicksort. Technical Report ALCOMFT-TR-03-50, ALCOM-FT, Barcelona, September 2003.

[174] Conrado Martínez and Xavier Molinero. An Efficient Generic Algorithm for the Generation of Unlabelled Cycles. Technical Report ALCOMFT-TR-03-5, ALCOM-FT, Barcelona, June 2003.

[175] Conrado Martínez and Xavier Molinero. Generic Algorithms for the Generation of Combinatorial Objects. Technical Report ALCOMFT-TR-03-4, ALCOM-FT, Barcelona, June 2003. Appearance: In B. Rovan and P. Vojtáš, editors, *Proceedings of the 28th Int. Symp. Mathematical Foundations of Computer Science (MFCS)*, volume 2747 of *Lect. Notes in Comp. Sci.*, pages 572–581. Springer-Verlag, 2003.

[176] Conrado Martínez, Daniel Panario, and Alfredo Viola. Adaptive Sampling for Quickselect. Technical Report ALCOMFT-TR-03-2, ALCOM-FT, Barcelona, April 2003. Appearance: In *Proc. of 15th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2004. To appear.

[177] Klaus Meer. On consistency and width notions for constraint programs with algebraic constraints. Technical Report ALCOMFT-TR-03-33, ALCOM-FT, Aarhus, September 2003. Appearance: *LNCS*, 2441:88–102, 2002.

[178] Klaus Meer. On the complexity of some problems in interval arithmetic. Technical Report ALCOMFT-TR-03-34, ALCOM-FT, Aarhus, September 2003. Appearance: *LNCS*, 2747:582–591, 2003.

[179] U. Meyer, P. Sanders, and J. F. Sibeyn, editors. *Algorithms for Memory Hierarchies*, volume 2625 of *LNCS*. Springer, 2003.

[180] Ulrich Meyer and Norbert Zeh. I/O Efficient Undirected Shortest Paths. Technical Report ALCOMFT-TR-03-61, ALCOM-FT, MPI, October 2003. Appearance: In *Proc. 11th Ann. European Symposium on Algorithms (ESA)*, LNCS. Springer, 2003, to appear.

[181] Friedhelm Meyer auf der Heide, Christian Schindelhauer, Klaus Volbert, and Matthias Grünewald. Congestion, Dilation, and Energy in Radio Networks. Technical Report ALCOMFT-TR-03-128, ALCOM-FT, Paderborn, December 2003. Appearance: In *ACM Symposium on Parallel Algorithms and Architectures (SPAA'02)*, pages 230–237, 2002.

[182] Peter Bro Miltersen, Jaikumar Radhakrishnan, and Ingo Wegener. On converting CNF to DNF. Technical Report ALCOMFT-TR-03-35, ALCOM-FT, Aarhus, November 2003. Appearance: In *Proc. MFCS'03*, Lecture Notes in Computer Science, 2003.

[183] Sotiris Nikoletseas, Ioannis Chatzigiannakis, Athanasios Antoniou, Haris Euthimiou, Athanasios Kinalis, and George Mylonas. Energy Efficient Protocols for Sensing Multiple Events in Smart Dust Networks. Technical Report ALCOMFT-TR-03-162, ALCOM-FT, Patras, December 2003. Appearance: In *Proceedings of the 37th ACM/IEEE Annual Simulation Symposium (ANSS)*, 2004.

[184] Sotiris Nikoletseas, Vicky Papadopoulou, and Paul Spirakis. Radiocoloring Graphs via the Probabilistic Method. Technical Report ALCOMFT-TR-03-158, ALCOM-FT, Patras, December 2003. Appearance: In *Proceedings of the Fourth Panhellenic Logic Symposium*, pages 135–140, July 2003.

[185] Sotiris Nikoletseas, Grigorios Prasinos, Paul Spirakis, and Christos Zaroliagis. Attack Propagation in Networks. Technical Report ALCOMFT-TR-03-153, ALCOM-FT, Patras, December 2003. Appearance: *Theory of Computing Systems*, 36:553–574, 2003.

[186] Sotiris Nikoletseas and Paul Spirakis. Distributed Algorithms for Representative Problems in Ad-hoc Mobile Environments: A Critical Survey. Technical Report ALCOMFT-TR-03-170, ALCOM-FT, Patras, December 2003. Appearance: In *Book Chapter, to appear by World Scientific*, 2003.

[187] Sotiris Nikoletseas and Paul Spirakis. How to certify in deterministic polynomial time the high chromatic number of instances of sparse random graphs. Technical Report ALCOMFT-TR-03-181, ALCOM-FT, Patras, December 2003.

[188] Jordi Petit. Combining Spectral Sequencing and Parallel Simulated Annealing for the MinLA Problem. Technical Report ALCOMFT-TR-03-83, ALCOM-FT, Barcelona, November 2003. Appearance: *Parallel Processing Letters*, 13(1):77–91, 2003.

[189] Harald Räcke. Minimizing Congestion in General Networks. Technical Report ALCOMFT-TR-03-130, ALCOM-FT, Paderborn, December 2003. Appearance: In *Proceedings of the 43rd IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 43–52, 2002.

[190] Harald Räcke, Christian Sohler, and Matthias Westermann. Online Scheduling for Sorting Buffers. Technical Report ALCOMFT-TR-03-93, ALCOM-FT, Paderborn, November 2003. Appearance: In *Algorithms - ESA'2002, Proc. 10th Annual European Symposium*, 2002.

[191] Mireille Régnier and Alain Denise. Rare Events and Conditional Events on Random Strings. Technical Report ALCOMFT-TR-03-98, ALCOM-FT, INRIA, November 2003.

[192] Stefan Rührup, Christian Schindelhauer, Klaus Volbert, and Matthias Grünewald. Performance of Distributed Algorithms for Topology Control in Wireless Networks. Technical Report ALCOMFT-TR-03-126, ALCOM-FT, Paderborn, December 2003. Appearance: In *17th International Parallel and Distributed Processing Symposium (IPDPS'03)*, 2003.

[193] Kay Salzwedel. Algorithmic Approaches for Storage Networks. Technical Report ALCOMFT-TR-03-115, ALCOM-FT, Paderborn, December 2003. Appearance: Springer-Verlag Inc., New York, NY, USA, 2003.

[194] Guido Schäfer and Naveen Sivadasan. Topology Matters: Smoothed Competitiveness of Metrical Task Systems. Technical Report ALCOMFT-TR-03-110, ALCOM-FT, MPI, November 2003. Appearance: In *Conference Proceedings of the 21st International Symposium on Theoretical Aspects of Computer Science (STACS)*, 2004. To appear.

[195] Stefan Schamberger. Heuristic Graph Bisection with Less Restrictive Balance Constraints. Technical Report ALCOMFT-TR-03-142, ALCOM-FT, Paderborn, December 2003. Appearance: In *International Conference on Parallel and Distributed Processing Techniques and Applications, PDPTA'04*, to appear.

[196] Stefan Schamberger. Improvements to the Helpful-Set Heuristic and a New Evaluation Scheme for Graphs-Partitioners. Technical Report ALCOMFT-TR-03-143, ALCOM-FT, Paderborn, December 2003. Appearance: In *International Conference on Computational Science and its Applications, ICCSA'03*, number 2667 in LNCS, pages 49–59, 2003.

[197] Stefan Schamberger and Jens-Michael Wierum. Graph Partitioning in Scientific Simulations: Multilevel Schemes vs. Space-Filling Curves. Technical Report ALCOMFT-TR-03-144, ALCOM-FT, Paderborn, December 2003. Appearance: In *Parallel Computing Technologies, PACT'03*, number 2763 in LNCS, pages 165–179, 2003.

[198] Christian Schindelhauer, Tamás Lukovszki, Stefan Rührup, and Klaus Volbert. Worst Case Mobility in Ad Hoc Networks. Technical Report ALCOMFT-TR-03-125, ALCOM-FT, Paderborn, December 2003. Appearance: In *Proc. of the 15th ACM Symposium on Parallel Algorithms and Architectures (SPAA'03)*, pages 230–239, 2003.

[199] Meinolf Sellmann and Torsten Fahle. Constraint Programming based Lagrangian Relaxation for the Automatic Recording Problem. Technical Report ALCOMFT-TR-03-174, ALCOM-FT, Paderborn, December 2003. Appearance: *Annals of Operations Research*, 118:17–33, 2003.

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

[201] Paul Spirakis and Christos Zaroliagis. Distributed Algorithm Engineering. Technical Report ALCOMFT-TR-02-141, ALCOM-FT, Patras, August 2002. Appearance: In *Experimental Algorithmics – From Algorithm Design to Robust and Efficient Software*, chapter 10, pages 197–228. Springer-Verlag, 2002.

[202] A. Stewart, M. M. Clint, and J. Gabarró. Barrier Synchronisation: axiomatisation and relaxation. Technical Report ALCOMFT-TR-03-41, ALCOM-FT, Barcelona, September 2003. Appearance: *Formal Aspects of Computing*, 2003. To appear.

[203] Alexander Tiskin. Communication-Efficient Parallel Gaussian Elimination. Technical Report ALCOMFT-TR-03-6, ALCOM-FT, Warwick, June 2003. Appearance: In V. Malyshkin, editor, *Proceedings of PaCT'03*, volume 2763 of *Lecture Notes in Computer Science*, pages 369–383. Springer-Verlag, 2003.

[204] Alexander Tiskin. Packing tripods: A computational approach. Technical Report ALCOMFT-TR-03-7, ALCOM-FT, Warwick, June 2003. APPEARANCE: In *Proceedings of EuroComb'03*, pages 353–357. Charles University, Prague, 2003.

[205] Alexander Tiskin. Packing tripods: narrowing the density gap. Technical Report ALCOMFT-TR-03-90, ALCOM-FT, Warwick, November 2003.

[206] Ann Vandevelde, Han Hoogeveen, Cor Hurkens, and Jan Karel Lenstra. Lower bounds for the head-body-tail problem on parallel machines: a computational study of the multiprocessor flow shop. Technical Report ALCOMFT-TR-03-192, ALCOM-FT, Utrecht, December 2003. APPEARANCE: This paper will appear in INFORMS Journal on Computing.

[207] Peter Verbaan, Jan van Leeuwen, and Jiří Wiedermann. Lineages of Automata. Technical Report ALCOMFT-TR-03-195, ALCOM-FT, Utrecht, December 2003.

[208] Klaus Volbert. A Simulation Environment for Ad Hoc Networks using Sector Subdivision. Technical Report ALCOMFT-TR-03-129, ALCOM-FT, Paderborn, December 2003. APPEARANCE: In *10th Euromicro Workshop on Parallel, Distributed and Network-ba sed Processing (PDP'02)*, pages 419–426, 2002.

[209] Thomas Wolle. A Framework for Network Reliability Problems on Graphs of Bounded Treewidth. Technical Report ALCOMFT-TR-03-193, ALCOM-FT, Utrecht, December 2003. APPEARANCE: In *13th International Symposium on Algorithms and Computation, ISAAC 2002, Proceedings*, pages 137–149. Springer, LNCS 2518, 2002.

[210] Martin Ziegler. Fast Relative Approximation of Potential Fields. Technical Report ALCOMFT-TR-03-88, ALCOM-FT, Paderborn, November 2003. APPEARANCE: In *Proc. 8th Workshop on Algorithms and Data Structures*, volume 2748 of *Springer LNCS*, 2003.