Final Report

ALCOM-FT

Algorithms and Complexity
Future Technologies

Project No. IST-1999-14186

January 2004

Summary

This is the final 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 entire project period from June 1, 2000 to November 30, 2003.

ALCOM-FT has brought together eleven of the leading groups in algorithms research in Europe in a project aiming at discovering new algorithmic concepts, identifying key algorithmic problems in applications, and contributing to the 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 focus of the project has been a combination of application oriented research in three important areas – *massive data sets*, *massive and complex communication*, and *complex problems in production and planning* – with innovative methodological work on *experimental algorithmics* and *generic algorithmic methods*.

The project has been divided into work packages, one for each of the five areas above, and one additional dealing with management and dissemination efforts. The scientific advances made by the project consortium in each of these areas have been very notable in terms of quality as well as quantity, and the goals set forth in the project description have been reached. Part I of this document gives an overview of these achievements.

All deliverables have been delivered[1], and the large majority of these were delivered on schedule. Table 2 lists all deliverables and their delivery times. Further descriptions of their contents are given in Part II of this document.

The largest deliverable is D1, which constitutes a total of 538 scientific reports produced during the project period. This high productivity has not been achieved at the cost of quality, as can be seen by the fact that at least 410 of the 538 reports have already been published in the scientific community via key conferences and journals, including the most prestigious in the world.

The set of deliverables also contains a number of software projects contributing to the transfer of advanced algorithmic methods into practice. These are: A library for external memory

[1]Except the Technology Implementation Plan, which is to be handed in by February 1, 2004

computation with a functionality similar to the widely used STL C++ library (D6, D15, D24); A demonstration package for methods in data mining (D16); A LEDA Extension Package for distributed algorithmic engineering, to facilitating simulation of large, hierarchical, or mobile networks and development of network optimization methods for resource allocation (D7, D17, D25); A BSP-style library which allows efficient implementation of BSP-style programs in dynamic distributed environments (D8, D18, D26); Software for production and transportation planning, along with a database containing real-life problem instances from the same area (D19, D27, D28); A completed software system (SCIL) for structured combinatorial optimization problems, including high-level constraints and their separation routines (D20, D23, D29); A set of tools for executing, documenting, visualizing, and reproducing algorithmic experiments (D5, D13, D21 D30).

Further dissemination efforts of the ALCOM-FT project include arrangement of eight summerschools, arrangement of several conferences and workshops (among these the 2000, 2001, and 2002 versions of ALGO, the premier algorithmic event in Europa), creation of a project web site (D2) and two web sites disseminating algorithmic knowledge (D11, D14), and industrial talks on algorithmic issues,

The project has been lead by a Consortium Board consisting of the scientific leader from each site. The following two adjustments was decided upon by the Board during the project: The algorithmics group at the University of Cyprus, headed by Prof. Marios Mavronicolas was added as an ALCOM-FT site for further strengthening of Work Packages 2 and 4, and the project was extended by six months to ensure proper finalization of some delayed deliverables (D14, D15, D19, and D21).

The rest of this document is organized into three parts. The first describes the progress achieved within the six work packages of the project. The second lists details of each deliverable of the project. The third lists all scientific reports published as part of the project.

No.	*Deliverable*	*Month Planned*	*Month Delivered*
D1	Research reports	Cont.	Cont.
D2	Project Presentation	3	3
D3	Dissemination and Use Plan	6	6
D4	Guidelines for algorithmic experiments (internal release)	6	12
D5	Testbed for experimental algorithmics (specification)	6	6
D6	External memory experimental platform (design)	12	12
D7	Distributed algorithmic engineering software package (prototype)	12	12
D8	BSP-style library for dynamic distributed environments (prototype)	12	12
D9	Production and transportation planning modeling report	12	12
D10	Guidelines for algorithmic experiments (external release)	12	36
D11	Algorithm Forum web site	12	12
D12	First progress report	12	12
D13	Testbed for experimental algorithmics (internal release)	18	18
D14	PR web-pages for algorithmics	18/36	36
D15	External memory experimental platform (prototype)	24/36	36
D16	Data mining demonstration package	24	24
D17	Distributed algorithmic engineering software package (beta version)	24	24
D18	BSP-style library for dynamic distributed environments (beta version)	24	24
D19	Production and transportation planning software prototype	24/36	36
D20	Description language for high level constraints in optimization	24	24
D21	Testbed for experimental algorithmics (external release)	24	26
D22	Second progress report	24	24
D23	Library of separation routines for ABACUS	30/36	36
D24	External memory experimental platform (final release)	36/42	42
D25	Distributed algorithmic engineering software package (final release)	36/42	42
D26	BSP-style library for dynamic distributed environments (final release)	36/42	42
D27	Production and transportation planning software user evaluation report	36/42	42
D28	Production and transportation planning problem instance database	36/42	42
D29	Complete software system for structured combinatorial optimization problems (SCIL)	36/42	42
D30	Testbed for experimental algorithmics (final release)	36/42	42
D31	Technology implementation plan	36/42	–
D32	Third progress report	36/42	42
D33	Final report	36/42	42

Table 2: The deliverables of the project. The notation 18/36 denotes the planned delivery month before and after the extension.

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: 80, of which 61 also appear in other work packages

Presentation

Traditional algorithm design is based on the von Neumann model, which assumes uniform memory access costs. However, our current state-of-the-art machines increasingly deviate from this model. From the registers, through very fast caches and a bit less fast internal memories, to increasingly complex secondary or even tertiary storage structures such as systems of parallel hard drives or robotized racks of CD-ROM units, algorithms that process large datasets need an awareness of where the accessed data is stored, and must take into account that the cost of memory accesses depends on the characteristics of the memory support. Typically, each level of this complex memory hierarchy, compared to its preceding level, has slower access time, larger capacity, and cheaper per-bit costs. For instance, an operation on registers in a modern but humble machine, operating at around a gigahertz as clock speed, takes a few cycles, whereas a memory access takes microseconds, or about 1000 times longer, and a disk access takes milliseconds, or one million times longer. These rough estimates give a hint of the problem faced, but actual computations with actual times of actual machines give even more daunting differences.

Even the constantly growing capacity of disk drives cannot satisfy the ever growing demand for flexible and efficient storage systems. Increasingly, commercial database systems offer support for including, as fields of the records in their tables, so-called BLOBS (Binary Large OBjectS) corresponding to e.g. images or music, which means single fields of typically 50Kb to 500Kb for a jpeg-encoded image or 1Mb to 4Mb for an mp3-encoded rock-and-roll song; and this, repeated over and over for all the tuples in the corresponding relation, and adding up to impressive quantities of memory usage.

A simple way to provide almost unlimited storage capacity is the use of many disk drives in parallel. RAID systems offer a handful of hard drives together, appropriately interconnected, but much bigger systems of parallel disk will be in use in the near future. The underlying architecture of such a *storage network* may vary significantly, and its design and use is far from trivial.

The smooth acceptance by the end user of such large storage systems depends on the employment of effective algorithmic techniques able to extract peak performances of such complicated systems without the user having to care about how. The record of the research in algorithms along the XX century proves that algorithmic work at different abstraction levels and target tasks has been able to meet the expectations of the user regarding computing power: we must do the same for virtually unlimited storage. Considering how Computer Science has coped with the situation in the past, we chose to organize our contribution to this future vision along three lines: the algorithms that have to be employed internally by the storage structures proper; algorithms that run on top of them to solve basic combinatorial problems, such as those found in our successful internal memory algorithmic library LEDA; and the algorithms themselves, on top of these, that solve user-oriented, industry-relevant or scientific tasks. Thus, the remainder of this report is structured in three sections, corresponding to these views.

A touchstone of the latter view is the analysis, not feasible up to now, of very large combinatorial structures, which we have chosen to perform as a testbed of our algorithms and their implementations. Along this line, we are also studying a number of problems that are popular in the Data Mining community, and believe that our results will contribute to the progress of this discipline, both in the layer of understanding and the layer of applications.

This report will neither present details of the contributions, nor include a discussion of the research that we feel yet as not so mature as to be very relevant (but which is necessary anyhow): see the yearly reports or the archive of technical reports for this. We concentrate instead on the most novel contributions and their significance.

Use of External Memory

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

From the first year of the project we studied the algorithmic support for parallel disks, where we have contributed innovative algorithms based on randomization, combined with redundant storage and sophisticated scheduling algorithms and a newly discovered duality between writing to disks and prefetching. Descriptions of the algorithms and their analysis are in [371], [496], [493], where a variety of practical situations such as data streams or variable-length blocks are considered.

However, we are aware that current storage systems lack the flexibility to adapt to a growing space demand. One way of meeting capacity and performance requirements is the usage of dedicated storage networks. These networks abstract from and generalize the parallel-disks model in that they consist of a number of arbitrary storage devices, connected by an arbitrary network. A number of requirements, like data distribution, balance of data accesses, heterogeneity, or adaptivity, have to be met to exploit the full potential of storage networks. In particular, whereas simple schemes are able to cope with homogeneous systems, where all the storage devices are essentially identical, and static, so that the full system is known in advance [122], we need solutions for the practical cases of heterogeneous (i.e. storage devices have different capacity and/or transfer rate) and adaptive (i.e. storage devices may enter or leave the system) settings. We developed two new placement schemes: SHARE and SIEVE [123]. Both of them are the first storage algorithms that permit an optimal adaptation to changing space requirements. In [491] we give an overview of existing techniques and introduce some more techniques for the efficient handling of many of these conditions simultaneously.

Altogether, we find ourselves in a position where it is realistic to pursue the implementation of a storage virtualization, where the storage units are extendable with little technological conditions and usable without the end user having to care too much about how: this has prompted us to suggest a new storage management environment, based on our new concepts, that is able to significantly reduce management costs and increases the performance and resource utilization of the given SAN infrastructure [121].

Basic Combinatorial Algorithmics

Cache-oblivious algorithms represent an excellent exponent of the approach to computer science research within the ALCOM consortium. Indeed, this work takes its starting position from the abstraction of technological issues, defines crisply a computation model so that formal analysis of the computational parameters of the algorithms is possible (though sometimes only feasible through strenuously hard mathematical work), and allows us to compare with more standard models and thus identify which aspects of the development have the potential for a practical impact; then we can concentrate on these and thus obtain deeply relevant contributions, such as the Lazy Funnelsort algorithm that, in most current machines, compares favorably in practice with the sorting algorithms that are usually considered best. Thus a sound, abstract, theoretically mathematical approach ends up providing an advance in the state-of-the-art practical systems.

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

Sorting is a basic task in information processing, and is also one of the few currently known tools for designing cache oblivious algorithms. Improving the practicality of sorting algorithms, as we do in [133], is therefore important, and, indeed, our advances provide cache oblivious algorithms for a number of problems in computational geometry.

A fundamental question is whether there is a *separation* in power between cache-oblivious algorithms and algorithms in the standard two-level model; we have contributed important results to this question. For instance, still with respect to sorting, we have proved that our algorithm Lazy Funnelsort is optimal, by proving a matching lower bound, and that therefore the cache-oblivious model is essentially different from the standard two-level model. This is one of the ways we have found of solving the separation question, and in a way that is practically important; other related contributions study permuting and searching problems [125], [83]. On the basis of the knowledge gained through this work, in [128] we describe how a fine tuning of our cache-oblivious algorithms can leverage the overhead imposed by the cache-obliviousness into a large number of advantageous characteristics, and still be competitive with the best practical algorithms for the same tasks. We see this contribution as, again, one more of the most influential work in the theory of algorithms.

Graph Algorithms

Hardly any other area of combinatorial algorithms is of wider practical use than graph algorithms. From layout and placement problems in the construction of the actual circuits that run our programs, through the identification of useless code by compilers, to facility location or operational transport and delivery problems in industrial logistics, the classical graph traversal problems breadth first search (BFS), depth first search (DFS) and single-source shortest-paths (SSSP) are more and more applied to larger and larger graphs.

Whereas BFS and DFS are easily solved sequentially using linear time on internal memory, fast and efficient external graph explorations were one of the big algorithmic challenges at the start of ALCOM-FT. The best external-memory algorithms then known for BFS and DFS required $\Omega(n)$ I/Os on general graphs. This was even true for *semi-external computing*, where it is assumed that node-related data can be stored internally. In [445] we gave the first external memory BFS algorithm which achieves $o(n)$ I/Os on arbitrary undirected graphs of bounded degree, or even of unbounded degree for planar embedded graphs [26], and extended the results to other problems such as DFS and SSSP. The idea of starting from suboptimal parallel algorithms gave an important progress for DFS [27]. Finally, as a landmark achievement, in [439] we presented the first external-memory BFS algorithm, fairly simple and likely to be practical, that achieves $o(n)$ I/Os on arbitrary undirected graphs.

A similar development was followed for SSSP problems. Various algorithms, focusing on the optimization of some or another aspect, were designed; some of them were obtained as follow-ups of algorithms for other traversals and appear in the previously cited reports, whereas others were described in [446], [448], [444], [447] and, finally, [449], which gives the best algorithm known for some specific applications (namely, for the ubiquitous case of sparse graphs with weights not too large, where our algorithm saves factors close to the square root of the block size). The specific advantages of each of the algorithms designed are briefly summarized in the corresponding yearly reports.

Applications

We have studied as well a number of algorithms and problems driven by more specific applications. Originated on the scientific interest in the combinatorial properties of graphs that appear in reality, such as call graphs or the WWW, and which are too large for treatment with internal memory algorithms, we have studied the power-law graphs that model them using our external-memory graph algorithms. We have found algorithms that, for the specific case of these graphs, became the first work-efficient $o(n^{1/4})$ average-case time SSSP algorithms [446]. In [514] we give a set of heuristics which allow to perform semi-external DFS for directed graphs in practice. The heuristics have been applied to models of the web graph and some call graphs, with impressive speed improvements, sometimes 100-fold better than any other alternative. Also, the topological properties of web-like graphs, including a real fragment of the WWW consisting of 200M HTML pages and 1.4 billion edges, have been studied thanks to our new algorithms for external or semi-external computation: [412] and [258]. This research has been facilitated by the `<STXXL>` library described below.

Approximation and Sampling

An important problem in the handling of large data sets is that of extracting relevant information from largely unstructured data, where the treatment often has to be based on statistical decisions. One then needs to be able to recognize what are statistically unavoidable regularities and distinguish these from the observed events that are likely to be meaningful. It is to be noted that classical probabilistic methods do not readily provide answers to these questions, while the original combination of formal and analytic approaches developed in the project enables us to attack the analysis of significance thresholds for an extremely wide class of complex patterns [306], [243].

A major new result is the LogLog Counting algorithm of [269] 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 with 1% or 2% of the actual value.

On relational data, a central problem in Data Mining applications is the computation of frequent sets from a given family of subsets of a fixed universe. This is actually a version of the problem of computing the transversal hypergraph of a given hypergraph, implicitly described by the frequencies of occurrence of items from the universe in the dataset, and is solved usually in practice by the so-called Apriori Algorithm. We proposed in [47] an alternative scheme based on a best-first strategy. Besides being efficient, it has the main advantage of combining naturally with adaptive sampling schemes, developed by our own researchers, which decide whether the sample size is sufficient for a given approximation quality, instead of using a batch sampling, i.e. a fixed sample size obtained from a standard Hoeffding-Chernoff analysis.

In fact, with respect to such (widely used) batch sampling, which is frequently inapplicable due to the extremely demanding sample sizes required, adaptive sampling maintains the same theoretically well-founded statistical guarantee of representativeness; simultaneously, in the vast majority of practical cases, it requires a reasonably small sample size. Therefore, as promised in the initial work plan, we have developed and programmed versions based on adaptive sampling of Data Mining and Machine Learning tasks, namely, the frequent sets problem as described above (where the standard schemes for frequent sets algorithms, like the Apriori method, are unable to take profit from the advantage of adaptive sampling, so that an alternative algorithm was developed), the construction of decision trees, and the construction of boosting-based weightings of decision stumps. These programs must be considered as research prototypes only and are accessible through our website. Finally, we should mention that we developed a general unified framework [347] to study such adaptive sampling schemes, ready for implementation and allowing also to prove limitations of the approach in the form of lower bounds.

With the additional ingredient of letting the underlying probability distribution, according to which the sampling is done, evolve in time, we have been able to apply sampling strategies to the training of Support Vector Machines (SVM) for classification problems; these algorithms are becoming another central topic in the datamining community. In [52] we provided an alternative randomized decomposition algorithm for which we could prove a strong bound on the expected completion time: it is the first (and so far the only) SVM training algorithm for which subquadratic expected running time bounds exist. An extension to regression problems appears in [53].

Software Libraries for External Memory

As planned, we have devoted considerable effort to making our most basic algorithms useful to persons outside the project. Initially we developed a single-disk external memory library [197] consisting of a low level block manager that allows the use of file systems and raw devices interchangeably and a higher level part that implements numerous algorithms and data structures such as sorting, matrix arithmetics, simple graph algorithms, priority queues, B-trees, buffer trees, and suffix arrays.

On the basis of this experience, and of wide previous experience in the design of algorithmic libraries, as explained in the initial proposal, we have implemented a first version of a platform for the experimentation with external memory algorithms, with a view to practical uses: our two main requirements were, therefore, very high performance and easy migration from existing software.

Based on an inexpensive Linux system that sustains full bandwidth I/O from up to eight disks, and using Posix threads and ordinary synchronous file system I/O calls, we could implement a novel I/O layer that supports both parallel disks and overlapping of I/O and computation. On top of this, a highly portable library is being built; by now it has support for almost all the STL of C++, while at the same time guarantees close to optimal use of all disks. Our implementation is two to three times faster than previous algorithm libraries, and is extremely innovative in that it is the first time that so many theoretical optimizations reach a fully practical, ready-to-use system. As a token of the outcomes of the experimentation made possible by this library, we can mention a new parallel disk sorting algorithm [226] that for the first time guarantees optimal overlapping of I/O and computation.

The part of this library relative to external memory graph algorithms is also growing fast; we already have an implementation for minimum spanning trees and connected components of massive graphs and there are ongoing implementation projects of external BFS, search trees, suffix tree construction, and other algorithms from [145, 386, 443]. More significantly, our library, called `<STXXL>`[2] has already suggested three possible cooperations outside our group, involving computer graphics, image processing, and geographic information systems. We see this interest of external potential cooperators as the best indicative of the appropriateness and timeliness of our whole developments.

WP 2: Networks and Communication

Participants: INRIA, MPI, Paderborn, CTI, Rome, Warwick
Coordinating site: CTI
Work Package Leader: Christos Kaklamanis
Number of technical reports: 216, of which 139 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 related to WP2 within the project has focused on two central issues:

- *Efficient communication in modern networks*
- *Distributed dynamic environments.*

In the rest of this section we overview important related results focusing on the theoretical developments. We cover the following specific research areas: Wavelength optimization in all-optical networks, algorithms for wireless networks, modelling of telecommunication networks, network routing, non-cooperative networks, graph partitioning and load balancing, scheduling problems, data management in networks, and BSP algorithms.

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

Wavelength optimization in all-optical networks

Our work in this area focused on *Wavelength Division Multiplexing* (WDM) all-optical networks. In such networks, communication requests are ordered pairs of nodes to be thought of as transmitter-receiver pairs and the WDM technology establishes connectivity by finding paths connecting transmitters to receivers and assigning wavelengths (or colors) to the paths so that paths sharing the same fiber are assigned different wavelengths. Since the number of wavelengths is limited, an important (and hard) optimization problem to be solved is to minimize the number of wavelengths used. This is known as the *wavelength routing* problem.

Within the project, we considered the wavelength routing in trees and proposed novel methods for wavelength optimization. These include the use of randomization and the exploitation of fractional path colorings for approximating the minimum number of wavelengths. In the first approach, we extended the notion of deterministic greedy algorithms to randomized ones, we studied their limitations, and we obtained new upper bounds for the problem which beat the known lower bounds for deterministic algorithms. For the analysis, we developed tail inequalities for hypergeometrical-like distributions. The relevant results are reported in [36, 163, 164].

In the second approach, we showed how to compute almost optimal fractional path colorings (these are relaxations of wavelength routing solutions) in bidirected trees of bounded degree, and we proposed new randomized rounding methods which, combined with known deterministic or randomized algorithms, yield solutions with substantially better approximation ratios. The second approach is general; it allows to improve the known approximability bounds for the wavelength routing problem for all networks where fractional path colorings can be computed efficiently. Furthermore, it can be applied in the case of multi-fiber WDM networks. In [154, 155], we apply this approach to (multi-fiber) bidirected trees and undirected rings. In the analysis, we use new tail bounds for interesting generalizations of classical occupancy problems. Other investigations related to wavelength optimization are reported in [152].

Algorithms for wireless networks

In the context of wireless networks, we have mainly considered ad hoc wireless networks (i.e., networks in which no infrastructure exists and communication is achieved by direct message exchanges between the nodes) with omnidirectional transmitters. We have studied issues related to their modelling, their dynamics and mobility, frequency assignment, and energy consumption.

Many reports were devoted to modelling of ad hoc networks. In [451], we studied the problem of path selection in static ad hoc networks on the two- and three-dimensional space. Measures of interest are the congestion, dilation and energy, while we took interferences between links into consideration. Our results provided interesting trade-offs between congestion, dilation and energy expressed in terms of the network size and the number of interfering links. We have also extended our investigations in the case of unidirectional transmission model where we allow communication in parallel sectors [253]. Other problems related to modelling and basic operations in mobile ad hoc networks were studied in [110, 179, 181, 183–185, 189, 465].

In particular, *smart dust* (or *sensor networks*) is a set of a vast number of ultra-small fully autonomous computing and communication devices, with very restricted energy and computing capabilities, that co-operate to quickly and efficiently accomplish a large sensing task. Smart dust can be very useful in practice, e.g., in the local detection of a remote crucial event and the propagation of data reporting its realization. The work [186] opened the basic algorithmic research on smart dust systems, by providing a simple and realistic model for smart dust, and by giving a set of simple and very efficient protocols for local detection and propagation along with a rigorous average case performance analysis. Subsequent studies are reported in [25, 111, 174, 175, 180, 188, 274, 458].

Concerning dynamics of ad hoc wireless networks, we compared the ability of several topologies to handle dynamic changes where radio stations appear and disappear [359]. For this, we measured the number of involved radio stations and presented distributed algorithms for repairing the network structure. It turned out that some of the models studied combined good performance in terms of interferences, energy and congestion. Additionally to this kind of dynamics, stations may move freely and unpredictably. In [505] we investigated distributed algorithms for ad hoc networks with moving stations in worst case scenarios (by defining models which restrict worst case mobility in a reasonable way). Our results showed that it is possible to maintain persistent routes with nice communication properties. Another recently studied [257] combinatorial notion, the *infection time* of graphs, may also be useful in estimating the information propagation time in mobile ad hoc networks.

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 labelling and radiocoloring) or independent set problems in graphs. In this context, in [21, 328–330, 460] we studied the *radiocoloring problem* (RCP). Among other results, we obtained complexity and algorithmic results for hierarchical and periodic planar graphs. Besides their inherent combinatorial interest, these graphs are motivated by large, structured wireless networks and can model periodic frequency assignment requests. New coloring algorithms for efficiently coloring squares of planar graphs were presented in [22–24]. We also studied another version of the problem in [97] where some nodes may have a pre-assigned frequency, and we obtained several algorithmic, combinatorial, and complexity-theoretic results.

We also considered online versions of the independent set problem in *cellular wireless networks* in [159–162]. These are planar networks with specific structure modelling interferences between base stations with equivalent transmitters uniformly distributed on the plane. We presented 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.

Energy is another scarce resource for ad hoc networks. In [78, 338] we studied 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 developed 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 [157, 158], we addressed the issue of supporting typical communication patterns like broadcasting, multicasting, and gossiping in wireless networks using the minimum amount of energy. We formulated a series of *wireless network design problems* which are shown to be NP-hard. Roughly speaking, these problems can be generally stated as computing an energy assignment to the nodes of the networks so that the supported connections maintain a connectivity property and the total energy is minimized. We considered the problems of minimum-energy broadcasting, multicasting, and gossiping in networks with symmetric and asymmetric edge-cost functions (representing required energy for communication between pairs of nodes). In asymmetric graphs, we obtained optimal logarithmic approximations for broadcasting and gossiping, and polylogarithmic inapproximability results for multicasting. For networks with symmetric edge-cost functions, we obtained logarithmic approximations for multicasting and constant approximations (and inapproximability results) for group communication problems. Our results exploit the relations of these wireless network design problems to node-weighted versions of famous combinatorial problems like directed/undirected steiner tree/forest, connected dominating set, and more.

Modelling of telecommunication networks

In this context, we aimed at investigating the basic features of mathematical models describing telecommunication networks. Our objective was 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 were analyzed through Markovian methods. The renormalization (Euler's scaling) was the main fundamental tool used to investigate the behavior of complex multi-dimensional Markov processes involved.

Our research work includes fundamental studies of the renormalized processes involved in the stochastic models describing telecommunication networks [214, 334, 335], studies of bandwidth allocation problems in TCP-like networks [265,267], investigation of the behavior of TCP congestion avoidance mechanisms [263, 360, 361] and admission control for Diffserv networks [264, 266, 377], stochastic modelling of TCP flows in ADSL traffic [82], and studies of the interactions of long and short TCP flows in IP networks [119, 362].

We have analyzed the *fluid limits* of Markov processes describing a telecommunication system. The fluid limits approach is a first order analysis of the state of complex systems and one of the most promising methods to tackle these difficult problems. Roughly speaking, a stochastic process is almost reduced to a deterministic dynamical system. Basic questions concern the relation between the initial process and the dynamical system. Our results showed that an old conjecture concerning the non-randomness of fluid limits in a classical Markovian model of a telecommunication network is wrong. Randomness implies that deterministic differential equations cannot describe completely the first order behavior of the network. We also showed that, in some specific cases, the limiting dynamical system involves products of random matrices. A complete characterization of such stochastic processes was achieved with the help of the dynamical system.

The study of AIMD (Additive Increase Multiplicative Decrease) algorithms is important since such algorithms are used in the congestion avoidance phase of TCP networks. An AIMD algorithm can be quickly described as follows: a source transmits W packets and, if an acknowledgement of reception is received, it then sends $W + 1$ packets (additive increase). Otherwise, a packet loss is detected and the source then sends δW packets (multiplicative decrease). The analysis focused on the evolution of the variable W (called the *congestion window size* for a very long connection: a source is sending a large file and the network losses packets from time to time. In our analysis we considered two models: a simplified one in which packet losses are independent (this model provides useful lower bounds on the behavior of AIMD algorithms), and another model where packet losses occur in groups. Our analysis of the second model showed that the simplified model underestimates the real performance of AIMD algorithms. For a fixed loss rate, correlation of packet losses improve the performance of the protocol (the higher the variability of the loss process, the better the throughput).

At the entrance of a network, admission control mechanisms determine whether new requests can be accepted or not. Connections are classified in several types according to their requirements (throughput, load, etc.) and the purpose of admission control is to accept the maximum number of connections and, simultaneously, to guarantee that each accepted connection will have a probability of packet loss less than ε (guaranteed quality of service). Algorithms based on a parameter called *effective bandwidth* are known to be very efficient if there is only *Assured Forwarding* traffic. In the context of Diffserv networks, there is also an additional class of traffic, the *Expedite Forwarding* traffic, which has absolute priority over every other traffic. Previous algorithms underestimate severely this high priority traffic. By considering a new stochastic model as a stochastic upper bound of the initial model, we proposed a new connection admission control algorithm. It uses effective bandwidths as well, but also considers the high priority traffic so that quality of service is guaranteed.

The aim of [82] was to characterize the Internet traffic in order to get realistic traffic models. After an extensive study of trace captures on an Internet backbone link connecting ADSL users, an approximation model was proposed and mathematical expressions for the stationary behavior of the system were obtained. The theoretical results were partially validated with the ADSL traces.

We also investigated the performance of a network link supporting long elastic 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. 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. We remark that it is generally difficult to get explicit results concerning queues with impatience. Under some assumptions, our estimations shed light on fairness issues concerning long flows.

Network routing

In the context of packet routing, we have produced several new results for fundamental routing problems. These problems have been intensively studied before because they form major building blocks for the efficient use of parallel computers. The results cover not only low level primitives but also more complex operations. Some of the results we obtained are remarkable since most important questions seemed to be largely closed for a decade. A simple randomized permutation routing algorithm for hypercube networks [534] closes a long standing gap between the best known algorithm and the trivial lower bound. A similar improvement for the problem of broadcasting a long message from one processor to all others is described in [498]. For gossiping (a natural generalization of broadcasting), powerful heuristics have been developed that can be applied to any specific network as well as a new algorithm for butterfly networks that breaks a long standing barrier for very large networks. Generalizations of these problems have been studied in [492]. Many of the algorithms designed work robustly on the emerging main-stream architecture of clusters of symmetric multiprocessors.

We also investigated the possibility or impossibility, and the corresponding costs, of devising *bufferless 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*). In [147,148], we studied 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, while direct routing is a special case of bufferless routing in which packets must be routed *without conflicts* along specific paths. For both problems, we derived algorithms which guarantee routing time within logarithmic factors of the obvious $\Omega(C + D)$ lower bound. Our results for hot-potato routing apply to trees, while the results on direct routing apply to a variety of commonly used network topologies such as trees, meshes, hypercubes and butterflies and are complemented by related hardness statements.

Adversarial Queueing Theory is a currently major arena in which to study the *worst-case* performance of routing and contention-resolution protocols for communication networks. The worst-case setting is much more natural and realistic than settings making probabilistic assumptions about the injection of packets into the network and their associated delays (with classical queuing theory being the prime example for the latter settings). Our research activity in this arena has attempted to classify protocols and networks according to the *stability* inherited to the network–the property that the number of packets in the network remains *bounded* at certain rates of injection. Within the project, we made substantial progress in several stability-related questions. We have shown the smallest injection rate that leads to instability when used on a certain FIFO network and the largest injection rate that guarantees stability when used on *any* FIFO network. We have shown that certain contention-resolution protocols that are stable in isolation are not so when composed together in a heterogeneous network with sufficiently large injection rate. We have addressed the question of how network structure (size, diameter, maximum node degree, minimum number of node-disjoint paths, network subgraphs) affects the stability properties of networks, for various greedy protocols. We have also investigated the effect of dynamically changing the capacities of a network on its stability and instability properties and have proved lower and upper bounds for the instability and stability thresholds, respectively, for several natural greedy protocols such as SIS (shortest-in-system) and LIS (longest-in-system). The related results are reported in [11, 19, 20, 246, 397–401] and push much forward the current front of research on stability issues in the context of Adversarial Queueing Theory.

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 so that the congestion 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 [88, 485], 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 approach is based on a hierarchical decomposition of the network with certain properties. We obtained existential and constructive results for such decompositions which yield competitive oblivious algorithms with polylogarithmic (in the size of the network) competitive ratios with respect to the network links.

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.

In [208] we proved that the worst-case coordination ratio (defined as the maximum possible value over all Nash equilibria of the social cost over the *social optimum* which represents the best achievable social cost) on links with possibly different speeds is $\Theta\left(\frac{\log m}{\log\log\log m}\right)$ generalizing a previous result of [434]. In [402], the notion of Nash equilibria was extended to that of approximate equilibria. We also presented the first study [203] of the price of selfish routing under general monotone families of cost functions and cost functions from Queueing Theory. In particular, a precise characterization of cost functions having bounded/unbounded coordination ratio was given. For example, cost functions that describe the expected latency in queueing systems have unbounded coordination ratio. We also focused on cost functions describing the behavior of Web servers that can open only a limited number of TCP connections. We concluded that web servers that are allowed to reject requests can guarantee a high quality of service for every individual request stream even under relatively high injection rates.

In [327], we presented the first complexity results concerning the KP model. We proved that it is hard to compute the maximum social cost for any given Nash equilibrium, or to compute the best or worst Nash equilibrium. Providing algorithms to compute a pure Nash equilibrium from a given schedule without increasing the social cost, we obtained a polynomial-time approximation scheme (PTAS) to approximate a best-case pure Nash equilibrium for the model of identical links [342] and related links [301], respectively.

In [342, 424], we have also given 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. For a survey on results in the KP-model we refer to [302].

In [423] we considered the Hybrid model (which combines in an interesting way features of the KP model and the *Wardrop* model), and we evaluated 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 presented 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 proved that, for the model of identical users and identical links, the worst-case Nash equilibrium is the fully mixed Nash equilibrium. Finally, we presented a comprehensive collection of (usually tight) constant (that is, independent of m and n) lower and upper bounds on the coordination ratio.

In [341] we considered 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 and for proving bounds on the coordination ratio under this model.

In a similar context [77], we studied 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 presented a complete characterization of the *complexity of computing optimal pricing strategies* and of *computing best and worst equilibria*. Basically, we showed 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.

Scheduling problems

We have addressed several scheduling problems focusing on their on-line variants, where information about tasks is not known in advance. In [393] we studied the problem of assigning independent tasks to uniformly related machines when only limited online information is provided to the scheduler. This is an elementary scheduling problem having many applications in areas where the communication is costly and a scheduler should require the minimum online information in order to achieve a fixed performance, such as online allocation of decentralized resources. The setting studied in [393] is a general framework whose special cases are the classical multiple-choice games for the assignment of independent tasks to identical machines. The latter case has been the subject of intensive research for the last decade. The problem is intriguing in the sense that the natural extensions of the greedy oblivious schedulers, which are known to achieve near-optimal performance in the case of identical machines, have been proved to perform quite poorly when distinct machine speeds come into play.

We presented a rather surprising lower bound stating that any oblivious scheduler that assigns an arbitrary number of tasks to n related machines would need $\Omega\left(\frac{\log n}{\log\log n}\right)$ polls of machine loads per task, in order to achieve a constant competitive ratio versus the optimum offline assignment of the same input sequence to these machines. On the other hand, we showed that the only missing information for an oblivious scheduler to perform asymptotically optimally is the amount of tasks to be inserted into the system by providing a new oblivious scheduler that uses only $O(\log\log n)$ polls, along with the additional information of the size of the input sequence, in order to achieve a constant competitive ratio. The nature of the proposed scheduler is quite different from the existing schedulers for identical machines (which are mainly greedy schedulers) and is based on the concept of giving higher priority to the least favourable machines. This oblivious scheduler was used as the core of an adaptive scheduler that does not require the knowledge of the input sequence and yet achieves similar performance.

In [74, 75] we considered 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 [74] we studied 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 is possible for a semi-clairvoyant algorithm to have constant competitive ratio with respect to average flow time on one single machine. We settled this open question by giving a semi-clairvoyant algorithm that is $O(1)$-competitive with respect to average flow time on one single machine. We also presented 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 we showed 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 [75] we introduce the notion of smoothed competitive analysis of online algorithms. Smoothed analysis has been proposed 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. A randomized version of the MLF was studied in [72] and was shown to perform well from a worst case analysis point of view.

A survey of recent work on on-line scheduling to minimize average response time or related metrics, on single and parallel machines is given in [76]. Other works reporting on related results can be found in [65, 69, 71, 219, 298, 393, 416].

Graph-partitioning and load balancing

Within the project, we extensively studied the k-partitioning problem. The task is to divide the set of vertices of a graph equally into a given number of parts while keeping the number of crossing edges between vertices belonging to different parts as small as possible. We refer to such a number of edges as the *cut size* of the partition. The special case of partitioning the graph into 2 parts is called a *bisection*, and the minimal cut size of all balanced bisections of a graph is called its *bisection width* (*k-section width* in the general case). Its calculation is NP-hard for arbitrary graphs and remains NP-hard for regular graphs. We obtained several new bounds on the bisection width of regular graphs [209, 244, 455]. Results for 3- and 4-regular graphs are of special interest because these are the lowest non-trivial degrees. Moreover, there are some direct applications of these results, e.g., upper bounds on the bisection width of 4-regular graphs have successfully been applied to the configuration of transputer systems. In [513] we considered the graph-partitioning problem by relaxing the balancing constraint. We presented new and generalized lower bounds for the k-section width within this model.

As we mentioned above, there are graph-partitioning problems in a wide range of applications. For example, when executing processes on parallel computer systems, inter-processor communication is a major bottleneck. One way to address this problem is to minimize the communication between processes that are mapped to different processors. This translates to the k-partitioning problem of the corresponding process graph, where k is the number of processors. In [280], we proved new relations between the structure and the eigenvalues of a graph and we presented a new method to get tighter lower bounds on the k-section width related to classical spectral bounds previously known. Our method makes use of the level structure defined by the k-section. We defined some global expansion property and proved that for graphs with the same k-section width the spectral lower bound increases with this global expansion. These results were enhanced and generalized in [281, 283].

Another research topic was the study of the relation between the properties of a graph and its suitability as interconnection network. One of the fundamental properties of a graph is the number of distinct eigenvalues of its adjacency or Laplacian matrix. Graphs with small spectra exhibit many symmetry properties and are well suited as interconnection topologies. In particular, load balancing can be done on such topologies in a small number of steps. In [278, 279], we constructed families of graphs with bounded or polylogarithmic maximum degree with a small number of distinct eigenvalues.

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

In [284], we generalized existing schemes in order to deal with heterogeneous networks. The generalized schemes may operate efficiently on networks where every processor can have arbitrary computing power. [285] concentrates on heterogeneous networks where the capacities of the edges varying. It is shown that the convergence rate of diffusion algorithms can be improved using edge-weighted graphs without deteriorating the flows quality. Moreover, in the same work, we considered common interconnection topologies and demonstrate, how optimal edge weights can be calculated for the *First and Second Order Diffusion Schemes*. Using theoretical analysis and practical experiments, we showed that improvements can be achieved on selected networks.

[425] investigates how to schedule a flow on a synchronous distributed network. In this way, we concentrated on the second phase of the two step approach described above. The goal is to use the minimum number κ of rounds to reach the balanced state. Although this flow scheduling problem trivially leads to a formulation as a linear program which is solvable (in a centralized way) in polynomial time, the best distributed scheduling strategy so far needs at most $O(\sqrt{n})$ times the minimum number of rounds. We have shown that every distributed scheduling strategy requires at least $\frac{3}{2}$ times the minimum number of rounds, and, furthermore, we presented a distributed algorithm for scheduling flows on tree networks, which requires at most two times the optimal number of rounds.

In [537] *evolving tree computations* on circulant (rings with chords) and related graphs are analyzed. In an evolving α-ary tree computation, a complete tree grows level by level, i.e., every leaf generates α new nodes that become the new leaves. The load-balancing task is to spread the new nodes on a network of processors in the moment they were created in such a way that the accumulated number of nodes per processor, i.e., its load, is as close as possible to the average number of nodes per processor. In [478], we studied the load-balancing problem with respect to a demand-driven parallelization of a ray tracing algorithm.

Data management in networks

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

A hierarchical bus network uses hierarchically, tree-like connected buses as a communication network. New communication technologies like SCI (Scalable Coherent Interface) make such networks very attractive, because they allow their easy construction and guarantee reasonable communication performance. Such networks can be modeled as tree networks: leaves correspond to processors, inner nodes to buses, edges to switches, and bandwidths of inner nodes and edges are related to bandwidths of buses and switches, respectively. For these hierarchical bus networks, we investigated in [450] the static data management problem with the goal to minimize the congestion, i.e., the maximum load of all edges and inner nodes induced by the read and write frequencies, divided by the bandwidth of the edge or inner node, respectively. We showed that the problem is

NP-hard and presented an algorithm that calculates a 7-approximate placement of all objects on a hierarchical bus network in polynomial time.

With the widespread use of commercial networks, as, e.g., the Internet, it is more and more important to consider commercial factors within data management strategies. The goal in previous work was to utilize the available resources, especially the bandwidth, as good as possible. In [404], we presented data management strategies for a model in which commercial cost instead of the communication cost is minimized, i.e., we considered a metric communication cost function and a storage cost function. We introduced new deterministic algorithms for the static data management problem on trees and arbitrary networks aiming at minimizing the total cost. This was the first analytic treatment of this problem which is NP-hard on arbitrary networks. Our main result was a combinatorial algorithm that calculates a constant factor approximation for arbitrary networks in polynomial time. Furthermore, we presented an algorithm that calculates an optimal placement of all objects on a tree in polynomial time.

An important aspect of dynamic data management is data tracking. Consider an arbitrary distributed network in which large numbers of objects are continuously being created, replicated, and destroyed. A basic problem arising in such an environment is that of organizing a distributed directory service for locating object copies. In [484], we presented a new data tracking scheme for locating nearby copies of objects in arbitrary distributed environments. Our tracking scheme supports efficient accesses to data objects while keeping the local memory overhead low. We also have shown that the new tracking scheme adapts well to dynamic changes in the network.

BSP algorithms

Within the project, we focused on two main activities in the context of bridging-models for parallel computing, in particular, in the investigation of the so-called BSP (bulk-synchronous parallel machines) model. On the one hand, a programming environment where one can actually execute BSP programs efficiently was implemented. On the other hand, BSP algorithms for combinatorial problems were developed and analyzed and the impact of the model used on the design of these algorithms was investigated.

In the context of the first activity, Paderborn implemented a platform-independent programming environment continuing our previous work. Virtual processors were introduced and it was realized that their execution can be stopped on the current processor and resumed at a different one. This is implemented on a Linux workstation cluster. A prototypical load-balancing mechanism was integrated into this system. The current release of the implementation of the Paderborn University BSP library can be inspected on the web page

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

In the context of the second activity, we developed algorithms for the broadcast problem [380] that obey locality constraints of the parallel machine where it is intended to be executed. The running time of this method was proved to be optimal for a wide range of BSP parameters and locality assumptions. An overview of further impacts of the used model on the design of efficient parallel algorithms was presented in [452] where we showed how the aspects “latency” and “critical block-size” influence the development of fast algorithms for solving the multisearch problem.

Also, an automatic configuration approach for implementing complex parallel BSP algorithms was investigated. [109] reports on the results of this approach. Here, a parallel algorithm is described by a sequence of instructions and of subproblems that have to be solved by other parallel algorithms called as subroutines, together with a mathematical description of its own running time. There may also be free algorithmic parameters that have an impact on the running time. As the running time of an algorithm depends on several machine parameters, the actual composition of the parallel program for an actual parallel machine given all these ingredients is a difficult task. Such a configuration system has been implemented using the Paderborn University BSP library. As an instructive example, we chose to implement a collection of sophisticated minimum spanning tree algorithms on two parallel machines, using the PUB library. The results of these implementations can be found in [109].

Other topics

Finally, within the project, there has also been work in some other important topics related to WP2. These include investigation of network fault-tolerance issues, Web-related problems and algorithms, distributed algorithms, networked data structures, properties of random network models, network design, algorithms for graph coloring problems, algorithms for dynamic networks, contention resolution protocols, and other fundamental network problems.

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: 34, of which 20 also appear in other work packages

The goal of WP 3 was to explore novel algorithmic issues in the area of production and transportation planning. Here we have focused 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.

The large fraction of WP3 reports that also appear in other work packages emphasizes the fact that we were aiming at a combination of practical and fundamental research. Almost all of the reports have appeared in conference proceedings like ESA, ALENEX, ICALP, and CP-AI-OR or have appeared in/been accepted for journals like INFORMS Journal on Computing, Math Programming, OR Letters, and Journal of Heuristics. Researchers participating in WP3 have been active in the program committees of many international conferences like ESA, IPCO, ICALP, COCOON, and CP-AI-OR.

Below we give an overview of the work done and results that have been established in WP3 in 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

Planning

The seat reservation problem deals with the question of deciding which seat to issue to a customer given that he can be accommodated. It is a typical example of a planning problem that is complicated by a lack of knowledge about the future: each customer has to be served if possible, although later it may turn out that he had better been ignored. We have studied two variants of this problem: the fully on line version, in which a specific seat has to be issued immediately ([43], [117]), and the partially on line version, in which passengers are told that they will get a, yet unspecified, seat when they get on the train. We have developed several algorithms with optimal competitive ratio for a number of scenarios of the first variant and an algorithm that solves the second problem to optimality.

Network planning

Within WP3 we have analyzed a number of problems involving network optimization (see [441], [511], [94], [89]). In the latter two reports we present and compare several metaheuristic algorithms, one of which is the rather new technique of Ant Colony Optimization; it turns out that this technique not only works well, but it is well suited for dealing with dynamic networks as well.

Rostering

A major component in WP3 is the study of complex crew rostering problems. We have looked at problems in the airline industry ([512], [295], [358]) and in home health care (Deliverable D9 and report [86]). These problems are highly difficult to solve in practice as the basic assignment problem is perturbed by difficult rules and regulations reflecting legal aspects, contracts, preferences, etc. Moreover, rules and regulations may vary significantly from one company to another making it thus evident to focus on general models that can incorporate various rule structures. In close cooperation with our industrial partners, we have dealt with problems concerning the integration of yield management and fleet assignment facilities and the integration of the airline fleet assignment and aircraft rotation building phases. These have been tackled by applying a combination of column generation and constraint programming and through local search, and integer programming; this has led to encouraging results.

In close co-operation with industrial partners we have developed a compact model for the home health care problem that is flexible enough to break down most real-world home health care problems of different characteristics. Especially, in [86] 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 are 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.

Scheduling

Our work on different kinds of machine scheduling problems has led to a large number of reports (see Deliverables 9, 27, 28 and reports [369], [370], [6], [288], [287], [216], [7], [4], [368], [5], [532]). The deliverables concern a scheduling problem originating from the pharmaceutical industry. We have come a lot closer to solving this problem, but the algorithms need to be tested in practice. Then we can also see the effect of the randomness of the processing times, which has been studied in [4].

Methodological research

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

Constraint programming and column generation are two very important techniques, which are used in many of the algorithms that we have presented to solve the real-life problems. In many applications we have to solve some subproblem then, which usually is some kind of constrained standard problem, like a shortest path problem or knapsack problem. We have developed several algorithms to solve these quickly, either to optimality of approximately (see [441], [506], [293], [220], [297]).

Moreover, many of the algorithms we have proposed are enumerative of nature, and thus depend on how well the search space can be reduced. We present several methods to reduce the search space in the reports [296], [293], [510], [511], [509], [294], [297]; these all are very effective.

Finally, we have developed the package SCIL (Symbolic Constraints in Integer Linear Programming), which enables the user to define high-level constraints, which often occur in the problems encountered in WP3 (see Deliverable D20).

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: 358, of which 178 also appear in other work packages

Of the 358 reports produced within the WP4 work package over the life of the ALCOM-FT project, at least 181 are in the published proceeding of conferences and workshops, at least 71 others appear in, or are about to appear in, journals, while another 7 are in other published forms such as book chapters.

Generic Methods have a central role in ALCOM–FT, providing new algorithms and analysis techniques, and supporting the more applied research within each of the other work packages. This Final Report will summarize the major themes addressed during the last $3\frac{1}{2}$ years with a few illustrative examples from different areas. More information appears in the three Annual Progress Reports and full details are to be found in the Report Series, the major deliverable of this work package, which is available on-line in our Project's web pages.

Analytic combinatorics and random structures

Randomized algorithms are of enormous importance in computer applications today. One strand of research within WP4 is the precise analysis of discrete models of randomness. This is essential in providing tools to predict the practical complexity of probabilistic algorithms. There has been a major change of emphasis in the design and analysis of algorithms from worst-case performance to an appropriate average-case "typical" behaviour. Further, there have been greater efforts to estimate run-times and resource bounds more precisely than with the traditional "big-O" analysis. This developing area of "random combinatorics" arising from algorithmic design may be seen as a largely European endeavour, and the work done in ALCOM–FT is recognised world-wide.

Finite-state and context-free models arising from finite Markov chains, formal verification and other sources are analysed systematically in [310]. Applications of this analysis to communication protocols and RNA structure are provided in [265, 376]. An important kind of question concerning large databases of texts or biological sequences concerns the significance or otherwise of the occurrence of some given pattern. This arises also in such contexts as intrusion detection and filtering large traces of concurrent events in computer systems. The analytic techniques to predict the number of occurrences of such patterns in sequences obeying rather general randomness properties are provided in [112, 306], and applications are considered in [243, 488].

Probabilistic tools are needed for the design and analysis of protocols for the reliable transmission of information in distributed communication networks. ALCOM–FT research has concerned such important features as routing strategies [246, 434], data management [404], load balancing [85], contention resolution (as in Ethernet) [351], stability [214], reservation policies [334], and tradeoffs between density and robustness in interconnection networks [307, 459, 464]. Formal models representing some naturally-arising graphs are established in [248, 475].

Some examples from this "random combinatorics" work are [56, 60] in which the "kernel method" yields closed-form formulae and asymptotics for some constrained random walks, with applications to uniform random generation problems from combinatorics and algorithmics. A further example is [500], which investigates overlapping file allocation problems for disks and uses the Lambert function to bound the degree of overlap.

Two major publications [311, 312] establish a basis for combinatorial enumeration of random discrete structures and a corresponding asymptotic theory. The approach is founded on formal generating functions and their complex-analytic properties. The objects of study, such as words, trees, graphs, and permutations, arise in diverse applications from statistical physics, computational biology, probability theory, and the analysis of algorithms.

The techniques are illustrated in [262,303,309,313,487], with examples investigating connectivity in random graphs, patterns in random texts, divide-and-conquer algorithms, and the expected complexity of balanced trees. Applications range from data mining to intrusion detection and the interpretation of genomic sequences.

GCD algorithms arise in many important practical applications and their characteristics have been extensively studied. New results in [50] give very a precise formulation of their distributional properties.

Optimization and on-line algorithms

Substantial contributions in this area have been made by several groups within ALCOM–FT during the period of the Project. These range from methodological studies of the measures used to evaluate on-line algorithms to very detailed analysis of particular situations.

The quality of on-line algorithms is often measured by comparing them with (off-line) algorithms which have all of the input data available at the start of their operation. The *competitive ratio* of an on-line algorithm is the worst case ratio of the performance of the on-line algorithm and the performance of an optimal off-line algorithm. Although the competitive ratio is a popular analytic tool, in many cases it does not differentiate satisfactorily between different on-line algorithms.

In [113], a new measure, the *relative worst order ratio*, is defined which allows a direct comparison between different on-line algorithms, instead of comparing each with an optimal off-line algorithm. This measure is applied to the paging problem in [115]. Commonly used deterministic algorithms, such as LRU (Least Recently Used) and FWF (Flush-When-Full), cannot be distinguished by their competitive ratios, both achieving the best possible ratio for deterministic algorithms. We introduce a novel algorithm and show that in relative worst order ratio terms this is strictly better than LRU, which is strictly better than FWF. A different approach to the paging problem, using locality of reference, provides a useful comparison of algorithms using the *fault rate* and gives results closer to reality than those obtained for the competitive ratio [10].

We made advances in the classic problem of on-line job scheduling, where a sequence of jobs arrives on-line to a number of identical machines. We investigated scheduling algorithms to minimize latency and optimize the "quality of service" as perceived by the end user, in the scheduling model where jobs can be pre-empted (i.e., suspended and resumed). We considered the metrics of average flow time and average stretch (where the *stretch* of a job is the ratio of its flow time to its processing time), as well as a rather general combination of the two. *Clairvoyancy* means that a new job's processing time is known on arrival, whereas in *semi-clairvoyancy* there is only approximate knowledge of this. Results obtained include upper and lower bounds on average stretch for randomized algorithms in the clairvoyant case [69], tight bounds for the randomized Multi-level Feedback algorithm (which is used in the Windows NT and the Unix operating systems) in the non-clairvoyant case [72], and a semi-clairvoyant algorithm that is $O(1)$-competitive with respect to average flow time [74]. Results for the more general objective functions appear in [73] and [315].

We have new results in another classic problem domain: on-line bin packing. In [42], we show that the competitive ratio of the well-known First-Fit algorithm can be improved by rejecting some items of input even when there is space for them. In further studies of bin-packing [116,286], we show that some natural algorithms achieve constant competitive ratios when restricted to sequences that can be packed completely by an off-line algorithm.

An interesting twist to on-line scheduling is to consider the effect of introducing a *sorting buffer*, a random-access buffer which can be used to rearrange the input sequence before processing. For example, in [486] we consider a service station with a sorting buffer with storage capacity k. The input items to be processed may have different attributes. The service station benefits from having consecutive items with the same attribute and seeks to use the buffer to minimize the number of attribute changes in the sequence it processes. The problem is motivated by applications in computer science and economics. Our main result is a deterministic strategy with a competitive ratio of $O(\log^2 k)$.

A fine example of the successful combination of Constraint Programming with Operations Research is given in [295, 512]. A new approach to the Crew Assignment Problem uses complex rule systems and shortest path algorithms to provide a powerful optimization algorithm.

Smoothed analysis

Several ALCOM sites are involved in this new approach to measuring on-line performance, seeking to explain how algorithms can 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 and to determine whether the worst instances are rare exceptions in the "complexity landscape" of the algorithm.

Smoothed competitive analysis of online algorithms was introduced in [75] to analyze the *Multi-Level Feedback* (MLF) algorithm. MLF tries to minimize the total flow time on a sequence of jobs in the non-clairvoyant case. The exponential worst-case performance of MLF contrasts dramatically with its good behaviour in practice. Our smoothed competitive analysis explains this contrast and gives the first average-case analysis of MLF.

Online problems, such as the paging problem, the static list-accessing problem and the k-server problem, can be modelled as *metrical task systems*. The *work function algorithm* (WFA) of Borodin, Linial and Saks has a tight competitive ratio of $2n - 1$ for metrical task systems of size n. We give in [501] a smoothed competitive analysis of WFA, showing it to be much better than its worst-case competitive ratio suggests. Other successes for our applications of smoothed analysis are to the knapsack problem [80], and to the *motion complexity* of systems of moving objects [210].

Applying an analogous approach to algorithms running on *discrete* data structures leads to the concept of *discrete smoothed analysis* [59]. Classical algorithms such as Quicksort and shortest-path algorithms in weighted graphs are analysed in models where the role of local averaging is played by *partial permutation* (of the instances) or *partial randomization* (of some bits of the instances).

Property testing

A decision problem is to verify whether a given object possesses a certain property or not. A *property testing algorithm* has just to distinguish between the case when a given object has the property and the case when the object is far away from any object having the property. Sometimes this can be done much more quickly than the corresponding (exact) decision problem, e.g., in time independent of the object description size.

The decision problem for proper ℓ-colouring of a hypergraph is known to be NP-hard, however in [205] we give a testing algorithm that can distinguish between the hypergraph having a proper ℓ-colouring and being "far" from having such a colouring in time independent of the size of the hypergraph.

Approximate data structures for sets of moving points are called *soft kinetic data structures* in [204]. A property testing algorithm can be used to check whether the data structure is "far from correct" and if so then it is repaired. Soft kinetic data structures are presented for sorted arrays, balanced search trees, heaps, range trees and Euclidean minimum spanning trees.

The efficiency of property testing algorithms depends on the queries available to the algorithm. It is shown in [206] how the introduction of range queries allows much faster testing algorithms for some properties concerned with convexity and clustering.

Approximation methods.

An important goal is to find fast approximation algorithms with constant worst-case ratio. We outline a few examples of progress in this direction, reported over the period of ALCOM–FT.

One graph-theoretic problem arising from applications such as transportation, warehousing and networking, bandwidth minimization and radio labelling, is the *k-clustering* problem. An efficient approximation algorithm making use of an advice structure, the *d-octopus* of the graph, is described in [316].

A polynomial-time algorithm for a constant-factor approximation of a partition of a sequence into a minimum number of increasing or decreasing subsequences, or of a partially ordered set into chains and antichains, is described in [317].

A polynomial-time constant-factor approximation to the *time-dependent orienteering problem* (a dual to the time-dependent traveling salesman) problem is obtained in [318].

A parameter affecting the performance of a communication network based on a given undirected graph G, is the oriented diameter of G (the smallest diameter of any strongly connected orientation of G). A linear time algorithm for chordal graphs is given in [319] which either finds that no strongly connected orientation exists, or approximates the oriented diameter to within a factor of about two. Unless P=NP, there can be no polynomial-time algorithm giving better than a $\frac{3}{2}$-approximation to the oriented diameter of a chordal graph.

Randomized algorithms

The problems of approximately counting combinatorial structures and of sampling the structures almost uniformly at random are closely related. A popular method for sampling is to simulate an appropriate Markov chain, the chief technical difficulty being to bound the *mixing time* of the chain (the time required before the stationary distribution is (nearly) reached). Such bounds are studied in a number of different areas: counting unlabelled combinatorial structures [352], modelling the iterated Prisoner's Dilemma [271], contingency tables [198], H-colourings of graphs [198,270], and 3-colourings of the planar grid [349].

Randomized algorithms are particularly important in many practical applications involving sorting and selection. Appropriate use of randomization is very successful in giving simple efficient algorithms which avoid their worst-case scenarios with high probability.

The Quickselect algorithm finds the m^{th} smallest element from n. Simple variants are investigated in [433], such as the *proportional-from-s* method, where at each stage the pivot is chosen from a random sample of size s as the element with the same relative rank as the element sought. Our analysis shows that even the simple proportional-from-3 algorithm is often better than the conventional median-of-3. We investigate in detail the strategy ν-find, which is a generalization of proportional-from-3. Experimental and theoretical results support the contention that ν-find should be the method of choice in practice.

When only the first m elements in the sorted order of n elements need to be determined, a simple variant of Quicksort gives a more efficient algorithm [429], which is shown to outperform the commonly-used partial-sort function in C++'s Standard Template Library.

The report [211] analyses an extension of the recent randomized primality test proposed by Grantham. Grantham's test has a smaller probability of error than the classic Miller-Rabin test, in the worst case. Our new algorithm gives an improvement in the average case, which is more relevant in many applications, such as the generation of random primes in cryptography. The worst-case performance is also improved, by using the efficient cubic residuosity test from [213].

Complexity bounds.

Tradeoffs and similarities

Determining relationships between different complexity measures is important in understanding the intrinsic complexity of computational problems.

In low-level circuit complexity the apparently very different notions of depth-restriction and width-restriction are shown in [367] to be closely related, by locating width-restricted classes within the hierarchy defined by constant depth circuits between $\mathbf{AC}^0$ and $\mathbf{ACC}^0$. Further work [366] led to an alternative characterisation of $\mathbf{ACC}^0$ in terms of constant-width polynomial-size planar circuits.

A classic example of a tradeoff between time and space arises from sorting. The time–space product required for sorting integers over a given domain is known to be $\Omega(n^2)$, even for non-comparison-based sorting. For times less than $\Theta(n \log n)$, no matching upper bound was known. This was improved to cover time down to $\Theta(n(\log \log n)^2)$ (and down to $\Theta(n \log \log n)$ with high probability using a randomized algorithm) in [470].

Another time-space trade-off, for deciding whether any pair elements are "close", was shown for branching programs in [473], extending a result by Ajtai.

In external memory models, where time is measured by the number of I/O operations (accesses to external memory), techniques for deriving I/O-space trade-offs from time-space trade-offs are presented in [404], giving examples from sorting and the "element distinctness" problem.

The complexity of proofs by resolution is often measured in terms of *resolution-width*. We show in [34] how to use pebble games to characterize resolution-width. Using this approach we can show that resolution-width is always at most the resolution-space, solving a well-known open problem.

Hard problems

Approximate counting is considerably easier than exact counting in the context of the polynomial hierarchy. The #P-hard problems of exact counting (introduced by Valiant) are as hard as the whole polynomial hierarchy, whereas approximate counting is not much harder than NP. However, little is known in detail about the complexity of approximate counting. The relation of *approximation-preserving reducibility* is used in [272] and begins to reveal some of the landscape of approximate counting problems. Between the expected equivalence classes which are the "hardest" and "easiest" problems in #P, there is an interesting natural class represented by the problem of counting the number of independent sets in a bipartite graph.

The continuing efforts to locate more precisely the phase transition from nearly-all-satisfiable to nearly-all-unsatisfiable for random 3-SAT formulae is continued in [382]. An improved upper bound of 4.571 is given for the critical clause-to-variable ratio.

When formulae are converted between CNF and DNF forms the size may increase "exponentially". A more precise quantification of this sheds some light on the complexity of satisfiability algorithms. Results relating the size of the increase to the numbers of variables and clauses is given in [453].

The NP-complete problem of finding the chromatic number of a graph is known to be also hard to approximate. An improved algorithm for solving this problem exactly, running in time $O(2.4023^n)$, is presented in [150].

Constraint satisfaction problems

The complexity of a constraint satisfaction problem can be greatly reduced by the imposition of restricted tree-width on the constraint graph. Previous results of this kind were limited to constraint programs over finite domains. In [437], we show that a similar restriction, *k-consistency* for tree-structured problems with algebraic constraints permits backtrack-free algorithms. Further width notions are studied which give backtrack-free solution algorithms for more general classes of constraint systems. In particular, some new classes of efficiently solvable 0-1-programming problems are derived.

Allen's *interval algebra* is concerned with binary relations such as "precedes or overlaps" on time intervals. It is used in artificial intelligence research for temporal reasoning about the time intervals associated with sets of events. In general, it is NP-complete to decide whether some collection of binary qualitative constraints on a set of intervals is satisfiable. A long-standing open question has been to classify the complexity of this decision problem according to the restricted types of allowed temporal constraints. This was solved in [406] and further extended in [405].

New data structures.

ALCOM–FT has been very productive of ingenious new data structures for a wide variety of algorithmic problems.

Dictionaries

Dynamic *dictionaries* support search, insert and delete operations. There are many well known data structures for them, such as binary search trees (BSTs, for short) and skip lists. BSTs perform poorly when they are skewed, so many notions of balanced BSTs have been devised. The method presented in [489] achieves logarithmic worst-case cost by using the sizes of the subtrees as balancing information. It is simple, easy to implement and fast in practice, provides worst-case guarantees, and allows for rank operations. Empirical evidence suggests that it should be considered among the first choices when constructing a BST.

Relaxed balance is the term used for search trees where the rebalancing has been uncoupled from the updating. This has advantages in some real-time applications and in shared-memory architectures. In [292], a search tree with relaxed balance having a height bound of $c \log n$ for any $c > 1$ is presented. The previous best bound had $c = 1.44$. In [410], B-trees with relaxed balance are enhanced to include group updates, and in [374] we present variants of B-trees with relaxed balance with improved space utilization.

In many dictionary applications, the access probability of any given key is not fixed, but decreases with the time since the last access to that key. In [290, 291], a *lazy update* scheme is supported by a new data structure, the *biased skip list* (BSL). The lazy update scheme gives optimal search and insert/delete times in terms of the bias of the input distribution, and is shown to have advantages in practice on real and synthetic data.

Another important technique for the dictionary problem is *hashing*. In [468, 469, 471] we have both theoretical and practical contributions: a characterization of optimality in terms of the number of memory locations accessed, a new idea for a deterministic algorithms, and a practically promising new hashing algorithm, *Cuckoo Hashing*.

At the extremes of speed versus storage, a randomized dictionary can be constructed so that each query only accesses a single word of memory [199] and gives the correct answer with a high probability. The storage required is just $O(n \log \frac{2u}{n})$ for a dictionary of n elements from a universe of size u.

String processing

For index structures on strings and sequences, both suffix arrays and suffix trees have been commonly used. Important advances have now been made in the efficient construction of suffix arrays. The linear-time algorithm given in [386] matches that previously known only for suffix trees, thus strengthening the appeal of the simpler, more space-efficient suffix arrays. The algorithm is also readily parallelizable to an efficient implementation on several varieties of parallel machine. A related algorithm [145], which is comparison-based however, constructs suffix arrays in $O(n \log n)$ time. It is space-efficient and competes well in practice with all previous algorithms.

An important problem in text compression and computational biology is to count the number of occurrences of some given pattern in a fixed string. It is shown in [131] that a string of length n can be preprocessed in time $O(n \log n)$, following which any pattern count can be found in time linear in the pattern length.

Other data structures

Many operating systems use a memory management system which allocates and deallocates memory blocks whose sizes are powers of two. The *buddy system* requires time which is logarithmic in the memory size. A new scheme is given in [132] which requires only constant time per operation.

Refined Buneman trees are a favoured way of generating evolutionary trees from bioinformatic data. The new algorithm for their construction given in [127] requires time only $O(n^3)$ and space

$O(n^2)$. This marked improvement over previous algorithms allows the investigation of refined Buneman trees for realistically sized data sets.

Graph and network algorithms

We mention just a few of the many results obtained during the life of ALCOM–FT in the area of graph algorithms.

New techniques, described in [227, 235], allow the construction of a deterministic algorithm to maintain explicitly the fully dynamic transitive closure of a graph as a Boolean matrix, supporting updates in $O(n^2)$ amortized time, and queries in unit worst-case time, or, in the case of deletions only, supporting updates in $O(n)$ amortized time. We also construct a randomized Monte Carlo algorithm which breaks through the $O(n^2)$ barrier on the single-operation complexity of fully dynamic transitive closure by supporting updates in $O(n^{1.58})$ and queries in $O(n^{0.58})$ worst-case time.

In [234, 237, 238], we have studied algorithms for maintaining all-pairs shortest paths in a dynamic graph. We use a completely new approach which yields a fully dynamic algorithm for general directed graphs on n vertices with real edge weights that supports any sequence of operations in $O(n^2 \log n)$ amortized time per update and unit worst-case time per distance query. This is the first algorithm that solves the dynamic all-pairs shortest paths problem in its full generality, and our bounds improve substantially over previous results. For more complex queries, such as "What is the shortest path from vertex x to vertex y avoiding a failed link (u, v)?", but for static graphs only, we have a compact data structure with $O(n^2 \log n)$ space which supports a query time of $O(\log n)$ [242].

In many applications, the graphs and networks which arise may have only small values of some parameter such as cutwidth, treewidth, or degree, and efficient algorithms can be designed to take advantage of this. Examples of our work in this area are linear-time algorithms when the cutwidth is small [521], and polynomial-time algorithms when the treewidth and degree are small [522], and efficient algorithms for computing the minimum dominating set and related problems on planar graphs [9].

Experiments on computing the treewidth of networks, with applications to decision support systems and telecommunication, are reported in [395]. The modular decomposition of a network can be used to efficiently compute the treewidth [104], and a variant of treewidth to measure the cost of a tree decomposition is introduced [101], where algorithms to compute tree decompositions of minimum cost are given.

Other work on graph algorithms includes a deterministic algorithm for maintaining a t-spanner of a general weighted undirected graph under a sequence of edge insertions and deletions [41], and simple combinatorial algorithms for the minimum feedback-arc set and minimum feedback-vertex set problems that achieve an approximation ratio bounded by the length of a longest simple cycle of the digraph [241].

Randomized algorithms to compute the bisection width of random cubic graphs with n vertices are presented in [245], giving an asymptotic upper bound of $0.174039n$. We also obtain an asymptotic lower bound of $1.325961n$ for the size of the max bisection. The tight analysis is based on the differential equation method.

In [140, 141] we perform a detailed study of the computational complexity of the problem of colouring a planar graph with the minimum number of colours such that each colour class avoids one or more forbidden graphs as subgraphs. Results include a complete answer for the case of a single forbidden connected subgraph, and many results on *subcolourings* (where each colour class consists of a disjoint union of cliques).

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: 93, of which 64 also appear in other work packages

The main goal of WP 5 was to strengthen the role of experimental techniques in algorithmic research in order to increase the practical impact of algorithmic results and to speed up their dissemination. We will begin with high level activity into this direction.

Rudolf Fleischer (MPII), Bernard Moret (U. New Mexico), and Erik Meineche Schmidt (Aarhus) organized a Dagstuhl Seminar on Experimental Algorithmics in September 2000. The result was deliverable D10, a volume [483] in the Springer LNCS series which possibly contains the most comprehensive collection of articles about algorithm engineering methodology. There were summer schools on algorithm engineering and algorithm library design organized by the Rome site and MPII in September 2001 and September 2003 respectively. The Workshop on Algorithm Engineering (WAE) was organized by MPII in 2000, by the Aarhus site in 2001, and became the "Engineering and Applications" track of the European Symposium on Algorithm in 2002 that was organized by the Rome site.

Of the 93 ALCOMFT reports produced in WP5, at least 76 have appeared in conference proceedings or journals, or are scheduled to do so soon. The large fraction of reports listed with at least one other work package underlinings the fact that algorithm engineering is an integral component of algorithmic research.

In the following we highlight some areas where we have worked very intensively and obtained interesting results that are easy to explain. It should be kept in mind however that this is an incomplete list. In particular, algorithm engineering permeates all work packages so that you will find several results in other sections of this report. For example, the work on cache efficient algorithms that was listed as part of WP 5 in the proposal turned out to have a larger theoretical component than anticipated and also significantly overlaps with WP 1 so that we want to refer to this section in particular.

Dictionary Data Structures

Dictionaries maintain a set of elements supporting fast search, insertion, and deletion of elements. Although from a theoretical perspective this area looked largely settled since the mid 1980s, an algorithm engineering perspective led us to a number of interesting results [83, 135, 225, 290, 331, 374, 471, 494] For example, priority queues [134] and search trees [83, 135] can also be implemented to run efficiently on all levels of the memory hierarchy simultaneously using the simple concept of cache oblivious algorithms. For more details refer to the description of WP 1.

We could also demonstrate that search trees that exploit integer keys [225] can achieve practical speedup over comparison based data structures.

Rasmus Pagh and Fleming Rodler received the best student paper award for a practical hash table data structure that has worst case constant access time [471]. A simple generalization of this *cuckoo hashing* approach additionally leads to highly space efficient hash tables [225].

Optimization Problems on Graphs

We have studied many optimization problems of graph that have many applications. For example the feedback arc set problem and its application to graph drawing [228] and local search algorithms for coloring and their application to timetabling and frequency assignment [22, 166, 167].

Here we report on three results in more detail that help to close gaps between theory and practice:

Previously, the best implementations of the maximum weighted matching problem uses suboptimal algorithms since the required data structures (mergable priority queues) for asymptotically better algorithms were considered too expensive. An implementation at MPII, that is now part of the LEDA library, shows that this is not the case — our asymptotically efficient algorithm outperforms previous implementation on almost all instances measured [440].

Steiner trees are one of the most widely studied network design problems with many applications ranging from VLSI design to bioinformatics. Reports [480–482] describe parts (new preprocessing and bounding heuristics) of a large program that can solve real world Steiner tree problems optimally although the problem is NP-hard. Our system is among the world leading systems in the sense that it is the fastest program for most instances from the well-established benchmark library SteinLib and that for some large instances it is the only program finding optimal solutions.

Many NP-hard problems can be solved efficiently for graphs with small treewidth. Although there are many theoretical result based on this observation, there were few practical implementations so far. One of the main problems was that it is nontrivial to find good tree decompositions. Reports [102, 103, 275] address this problem. Tests on networks derived from applications like decision support systems and frequency assignment have shown that the methods are usually very effective.

Dynamic Graph Algorithmics

In the last 15 years, many interesting theoretical results have been obtained on the practically interesting question how to maintain solutions to fundamental graph problems when only a few edges change. However, there have been few implementations and it was not always clear whether simple minded solutions would not outperform sophisticated worst case efficient algorithms for practical inputs. We have therefore conducted several experimental studies on transitive closure [337], single-source shortest paths [233], all pairs shortest path [237, 239], and minimum spanning trees [172]. A detailed survey was prepared in [539]. 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

Our most influential contribution to algorithm visualization so far is Leonardo, an integrated environment for visualizing C programs [229]. Leonardo is distributed with a collection of more than 60 animations of algorithms and data structures including approximation, combinatorial optimization, computational geometry, on-line and dynamic algorithms. Leonardo has been widely disseminated on CD-ROM in computer magazines and is available for download in several software archives over the Web. It has received several technical reviews and more than 18,000 downloads over the last two years. A newer extension is Leonardo Web, a tool for building animated presentations that also work over the Internet.

We are also studying advanced aspects of algorithm visualization. CATAI (Concurrent Algorithms and data Types Animation over the Internet) [173] allows multiple users to watch a visualization and interact.

In [108] 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.

Interactive Rendering of Highly Complex Scenes

In [389, 536] we present a new rendering algorithm. The *randomized z-buffer algorithm* uses a sampling technique to achieve interactive frame rates even for huge scenes with up to 10^{14} (implicitly defined) triangles. In contrast, traditional z-buffer algorithms have a running time proportional to the scene size and are thus unable to process huge scenes. Our system is also able to handle scenes that are so large that they can only be stored on a local or remote hard disk.

Report [479] presents a prototype of an e-commerce system that simulates realistic lighting of large scenes on high performance parallel computers.

Uniform Random Generation of Combinatorial Objects

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

The article [261] explains how to generate combinatorial objects of size between n and $n*(1+\epsilon)$ in linear time using a rejection method. In contrast, the best known algorithms that produce objects of size exactly n take quadratic time.

The report [262] proposes a new approach to the structure of random objects like constrained sequences, trees, graphs, or permutations, based on the idea of Boltzmann models from statistical physics. Most of these algorithms are of subquadratic or even linear complexity and can be automatically generated from specifications. Implementations indicate that objects of size 100.000 can now be generated where only size 100 was previously possible.

Engineering Distributed Algorithms

Mobile ad-hoc networks [176, 178, 181, 183, 184] and sensor networks [25, 174, 175, 180] are a challenging new area for distributed algorithms. Since these models are very complex, simulations of new algorithms are an essential part of algorithmic research in this area. Our approach is based on a lightweight simulation environment using C++ and LEDA, which allows us to make extensive simulations runs involving thousands of nodes.

We have also studied attacks in computer networks [461] and routing torus-connected processor arrays.

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

Overall, the management of the project has run very smoothly, and few problems have been encountered. Scientific productivity has been high, work has progressed as expected, funds have been used in good accordance with the budget, and communication within the consortium, as well as between the consortium and the EC, has worked well.

The Consortium Board, consisting of the scientific leader of each of the eleven participating sites and chaired by the Coordinator, has carried the overall responsibility for the management of the project. 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 have also conducted the reporting on the work done.

The Board has met at least once a year, monitoring the progress of the work and the use of the financial funds. Among the key decisions taken by the Board during the project is the inclusion of the algorithms group at the University of Cyprus, headed by Prof. Marios Mavronicolas, and the application to the EC for a six month extension in order to ensure proper finalization of some delayed deliverables (D14, D15, D19, and D21). Both steps proved successful: Site Cyprus has been well integrated in the consortium, as witnessed by the amount of cooperation between site

Cyprus and other ALCOM sites, and all deliverables have been delivered according to their new schedule.

Each year has featured an ALCOM-FT review meeting with participation of the EC project officer and review team. In connection with each review meeting, a consortium-wide scientific workshop has taken place, featuring talks on select work from each site.

Dissemination

The principal means of dissemination for the ALCOM-FT project is scientific reports describing the research done. In total, 538 scientific reports have been produced as part of the ALCOM-FT project. At least 410 of these[3] have already been published in the scientific community via key conferences and journals, including the most prestigious in the world. All 538 reports have also appeared in the ALCOM-FT Technical Report Series, which is available online.

A total of eight summerschools/workshops aimed at Ph.D.-students and young researchers have been organized using project funds, all with contemporary subjects in areas related to ALCOM-FT. The events were: *Advanced Course on the Foundations of Computer Science* (Saarbrücken, September 2000), *ALCOM Spring School on Dynamic Algorithms* (Paderborn, May 2001), *2nd Max-Planck Advanced Course on the Foundations of Computer Science* (Saarbrücken, September 2001), *ALCOM-FT Summer School on Algorithm Engineering* (Rome, September 2002). *Models and Algorithms for the Web* (Udine, 2002) *3nd Max-Planck Advanced Course on the Foundations of Computer Science* (Saarbrücken, September 2002) *Workshop on Parallelism in Algorithms and Architectures* (Paderborn, March 2003) *4th Max-Planck Advanced Course on the Foundations of Computer Science*(Saarbrücken, September 2003). All events had a high attendance, with participants from many European countries as well as from outside Europe.

Members of the ALCOM-FT project organized several internationally renowned conferences in areas highly relevant for the work of the project. These include ESA'00, WAE'00, APPROX'00 (site MPI), STOC'01, ICALP'01, SPAA'01 (site CTI), ESA'01, WAE'01, WABI'01 (site Aarhus), ESA02, APPROX'02 and WABI'02 (site Rome), CCC'03 (site Aarhus). Additionally, many smaller conferences and workshops were arranged by various ALCOM-FT sites, including two Dagstuhl meetings on algorithmic engineering.

An ALCOM-FT website (`http://www.brics.dk/ALCOM-FT`) has been set up, containing a description of the objectives of the project, the sites participating, and the deliverables produced. The website also contains the online version of the ALCOM-FT Technical Report Series. Several further deliverables have separate websites describing the work performed and reachable through the ALCOM-FT website. Two of these websites are directly concerned with dissemination, namely a site with PR-pages for algorithmics (among other things reporting success stories on collaboration among theorists and practitioners), and a website facilitating the sharing of lecture notes among lecturers within computer science.

To increase the awareness of algorithmics in research communities outside ALCOM-FT, we have also set up an algorithmic competition based on finding worst case instances for graph algorithms from the LEDA library (see `http://www.mpi-sb.mpg.de/~schaefer/MLLB/index.html`. Unfortunately, the competition has failed due to lack of entries submitted. Apparently, this venue of dissemination is not efficient, at least not in the form attempted here.

Dissemination of algorithmic awareness has also been fostered by talks given by partners to audiences in industry. In particular, the Paderborn site has established a forum called *Industrie trifft Informatik* (Industry meets Computer Science), where university and business cultivate their dialog in the field of information technologies and where possibilities for new cooperations are launched.

The consortium has maintained a high profile in Ph.D. education, illustrated by the fact that during the project, 47 persons obtained an ALCOM-FT related Ph.D.-degree.

[3]Based on the reporting from the authors. This is a lower bound, as it is quite possible that some publications have not been reported.

Evaluation

The ALCOM-FT project has been very successful. All deliverables have been completed. Scientific productivity has been high, and a substantial number of Ph.D.'s have been educated. The international standing of ALCOM-FT is illustrated by the large fraction of research reports which have been published, or been accepted for publication, at the most prestigious conferences in the field.

Part II Deliverable Overview

Below, we list the deliverables of the ALCOM-FT project grouped by work package. Deliverable D1 (538 research reports delivered continuously across the entire project period) is not included, as it pertains to all work packages.

The notation 24/36 denotes the planned delivery month before and after the six month extension of the project.

Deliverables of WP1

External Memory Experimental Platform

No.	*Deliverable*	*Month Planned*	*Month Delivered*
D6	External memory experimental platform (design)	12	12
D15	External memory experimental platform (prototype)	24/36	36
D24	External memory experimental platform (final release)	36/42	42

Description: Deliverables D6, D15, and D24 constitute the development of the `<STXXL>` library. The core of `<STXXL>` is an implementation of the well-known C++ standard template library STL for external memory computations, i.e., `<STXXL>` implements containers and algorithms which can process huge volumes of data that only fit on disks. The `<STXXL>` library is already used by several projects. Although still evolving, `<STXXL>` is by now a substantial piece of software with 14.000 lines of code and a number of statements comparable to the standard STL library. More information can be found at the `<STXXL>`-homepage at `www.mpi-sb.mpg.de/~rdementi/stxxl.html`.

Sites involved: MPI

Data Mining Demonstration Package

No.	*Deliverable*	*Month Planned*	*Month Delivered*
D16	Data mining demonstration package	24	24

Description: Deliverable D16 consists of programs that use our adaptive sampling techniques to scale up methods from Machine Learning and Data Mining. It contains four stand-alone programs (our algorithm for frequent sets as per [49], the scaled-up version of that algorithm as per [347], a program to experimentally evaluate association rules as per [170], and the scaled-up versions of the decision tree and decision stumps algorithms as per [347]) and two extensions (scaled-up weighted majorities of decision stumps and scaled-up K-Means) of Weka, a well-known and widely used open source data mining package. All the programs can be retrieved on the web from `http://www.lsi.upc.es/~gcasas/alcom/alcom_software.html`.

Sites involved: Barcelona

Deliverables of WP2

Distributed Algorithmic Engineering Software Package

No.	*Deliverable*	*Month Planned*	*Month Delivered*
D7	Distributed algorithmic engineering software package (prototype)	12	12
D17	Distributed algorithmic engineering software package (beta version)	24	24
D25	Distributed algorithmic engineering software package (final release)	36/42	42

Description: Deliverables D7, D17, and D25 are the Distributed Algorithms Platform (DAP). This is a software platform supporting the implementation, simulation, and testing of distributed algorithms. Its focus is on implementing distributed algorithmic ideas developed for wired and mobile wireless network systems. It is implemented in C++ using LEDA. To transfer the full power of LEDA to distributed experiments and implementation of distributed algorithms, DAP is provided as a LEDA Extension Package. The main strengths of DAP are:

- The user develops in a standard programming language (C++) and the programs can be easily and efficiently ported to a real distributed environment.
- A modular model is used, powerful enough to represent diverse distributed environments, yet simple enough to hide unnecessary details.
- A user is able to control, monitor, and visualize the simulation of his/her algorithm through a Graphical User Interface (GUI).
- Many users can simultaneously monitor and visualize (through a lighter version of the GUI) the simulation of some algorithm from remote locations.
- The execution of the simulation can be distributed among several hosts (if required).
- Randomness can be introduced in all areas of the simulated distributed system.
- Scenarios can be defined for repetitively simulation of particular situations.

More information can be found at the homepage of DAP at `ru1.cti.gr/~LEP-DAP/`.

Sites involved: CTI

BSP-Style Library for Dynamic Distributed Environments

No.	*Deliverable*	*Month Planned*	*Month Delivered*
D8	BSP-style library for dynamic distributed environments (prototype)	12	12
D18	BSP-style library for dynamic distributed environments (beta version)	24	24
D26	BSP-style library for dynamic distributed environments (final release)	36/42	42

Description: Deliverables D8, D18, and D26 constitute a contribution of the ALCOM-FT project to the further development of the Paderborn University BSP software library (PUB) for parallel computing. The PUB library is a BSP-style library that allows the efficient implementation of BSP-style parallel programs in dynamic distributed environments. Specifically, the deliverables embody work resulting in stable process migration procedures, load balancing routines, and the possibility for realizing any customized load-balancing methods. The ALCOM-FT funded final release 8.0 of the PUB library can be downloaded from `www.upb.de/~pub/` together with a comprehensive documentation.

Sites involved: Paderborn

Deliverables of WP3

Production and Transportation Planning Modeling Report

No.	*Deliverable*	*Month Planned*	*Month Delivered*
D9	Production and transportation planning modeling report	12	12

Description: D9 are two reports describing ways to model and solve optimization problems from two complicated real-life domains, namely scheduling multipurpose batch processes and rostering in home health care.

Sites involved: Paderborn, Utrecht

Production and Transportation Planning Software

No.	*Deliverable*	*Month Planned*	*Month Delivered*
D19	Production and transportation planning software prototype	24/36	36
D27	Production and transportation planning software user evaluation report	36/42	42

Description: D19 and D27 are two steps in the development of software for solving job scheduling problems with no-wait constraints.

Sites involved: Utrecht

Production and Transportation Planning Problem Instance Database

No.	*Deliverable*	*Month Planned*	*Month Delivered*
D28	Production and transportation planning problem instance database	36/42	42

Description: Deliverable D28 is a database of real life instances of the job shop scheduling problem. The instances stem from the pharmaceutical industry, and are here made accessible for research purposes. Because these instances come from an actual industrial company and may provide insight into their internal production methods, we cannot publish the described material on the web. If you are interested in this material, please write an e-mail to Peter Lennartz (`peterl@cs.uu.nl`).

Sites involved: Utrecht

Deliverables of WP4

Software System for Structured Combinatorial Optimization Problems

No.	*Deliverable*	*Month Planned*	*Month Delivered*
D20	Description language for high level constraints in optimization	24	24
D23	Library of separation routines for ABACUS	30/36	36
D29	Complete software system for structured combinatorial optimization problems (SCIL)	36/42	42

Description: Deliverable D29 is the completion of a software system for Symbolic Constraints in Integer Linear programming (SCIL). The system consists of a modeling language (developed as deliverable D20) together with standard symbolic constraints (developed as deliverable D23), the documentation, and examples. The system is now available as an easy to install software system which can be downloaded at `www.mpi-sb.mpg.de/SCIL/scil.tar.gz`. Further information can be found at the SCIL homepage at `http://www.mpi-sb.mpg.de/SCIL/`.

Sites involved: Cologne and MPI

Deliverables of WP5

Guidelines for Algorithmic Experiments

No.	*Deliverable*	*Month Planned*	*Month Delivered*
D4	Guidelines for algorithmic experiments (internal release)	6	12
D10	Guidelines for algorithmic experiments (external release)	12	36

Description: Deliverable D10 is a volume [483] in the Springer LNCS series, and is the outcome of a Dagstuhl Seminar on Experimental Algorithmics arranged by Rudolf Fleischer (MPII), Bernard Moret (U. of New Mexico), and Erik Meineche-Schmidt (Aarhus) in September 2000. It contains a comprehensive collection of articles on algorithm engineering methodology and algorithm engineering cases studies. Deliverable D4 was the table of contents of the (at that time) still unpublished volume.

Sites involved: Aarhus, MPI

No.	*Deliverable*	*Month Planned*	*Month Delivered*
D5	Testbed for experimental algorithmics (specification)	6	6
D13	Testbed for experimental algorithmics (internal release)	18	18
D21	Testbed for experimental algorithmics (external release)	24	26
D30	Testbed for experimental algorithmics (final release)	36/42	42

Description: Deliverable D5, D13, D21, and D30 constitute ExpLab, a set of software tools for executing, documenting, visualizing, and reproducing algorithmic experiments. The main goals of this tool set are:

- To provide a simple way to set up and run computational experiments.
- To provide a means of automatically documenting the environment in which an experiment is run so the experiment can be easily rerun (provided the same environment is still available) and the results can be more accurately compared to the results of other computational experiments.
- To eliminate some of the tedium involved in collecting and analyzing output by providing basic text output processing tools.

In particular, the point is not to replace existing tools that already provide useful functionality for computational experiments (such as gnuplot, make, perl, python). Rather, ExpLap extends the functionality already available, and provides new tools giving a more comfortable and productive experimentation environment with more easily reproducible results, thereby promoting scientific integrity. The software can be downloaded from ExpLaps web site `sourceforge.net/projects/explab/`.

Sites involved: MPI

Deliverables of WP6

Progress Reports

No.	*Deliverable*	*Month Planned*	*Month Delivered*
D12	First progress report	12	12
D22	Second progress report	24	24
D32	Third progress report	36/42	42
D33	Final report	36/42	42

Description: Deliverables D12, D22, D32, and D33 are the three progress reports and the final report for the ALCOM-FT project. They describe the work performed during the three periods June 1, 2000 to May 31, 2001, June 1, 2001 to May 31, 2002, and June 1, 2002 to November 30, 2003, as well as the entire project period June 1, 2000 to November 30, 2003.

Sites involved: All sites.

Project Presentation

No.	*Deliverable*	*Month Planned*	*Month Delivered*
D2	Project Presentation	3	3

Description: Deliverable D2 is the main ALCOM-FT web site, which contains a description of the project, a list of the participants, online access to all deliverables of the project, and information on further activities like summerschools. The web site can be found at `http://www.brics.dk/ALCOM-FT/`.

Sites involved: Aarhus

Web Sites for Dissemination of Algorithmic Knowledge

No.	*Deliverable*	*Month Planned*	*Month Delivered*
D11	Algorithm Forum web site	12	12
D14	PR web-pages for algorithmics	18/36	36

Description: Two websites launched as part of the dissemination effort of the ALCOM-FT project. Deliverable D11 is a site facilitating the sharing of lecture notes produced by members of the scientific community (inside as well as outside ALCOM). It can be found at `http://www.brics.dk/ALCOM-FT/LNC/`. D14 is a site aiming at describing algorithmic problems and solutions in a way which is easily understood by non-scientists, e.g. people from industry. It does so by reporting success stories on collaboration among theorists and practitioners, and by providing a repository of descriptions of problems and techniques. It can be found at `http://www.mpi-sb.mpg.de/~spyros/pr-www/`.

Sites involved: Aarhus and MPI

Dissemination, Use, and Technology Implementation Plans

No.	*Deliverable*	*Month Planned*	*Month Delivered*
D3	Dissemination and Use Plan	6	6
D31	Technology implementation plan	36/42	–

Description: Two required deliverables for IST projects. The documents delineate the consortiums plans for dissemination, use, and technology transfer of the results achieved in the project. The Technology Implementation Plan is due February 1, 2004.

Sites involved: All sites

Part III Scientific Reports

In the following, we list all 538 scientific reports (and two further ALCOM-FT related references [443, 483]) which have been published in the ALCOM-FT Technical Report Series. 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] M. Adler, F. Fich, L. A. Goldberg, and M. Paterson. Tight Size Bounds for Packet Headers in Narrow Meshes. Technical Report ALCOMFT-TR-01-9, ALCOM-FT, Warwick, January 2001. APPEARANCE: In U. Montanari, J. Rolin, and E. Welzl, editors, *Proceedings of ICALP 2000*, volume 1853 of *Lecture Notes in Computer Science*, pages 756–767. Springer-Verlag, 2000.

[3] Micah Adler, Harald Räcke, Naveen Sivadasan, Christian Sohler, and Berthold Vöcking. Randomized Pursuit-Evasion in Graphs. Technical Report ALCOMFT-TR-02-35, ALCOM-FT, Paderborn and MPI, May 2002. APPEARANCE: *Combinatorics, Probability and Computing*, 12(3):225–244, 2003.

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

[5] 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.

[6] Marjan van den Akker, Han Hoogeveen, and Steef van de Velde. Combining column generation and Lagrangean relaxation to solve a single-machine common due date problem. Technical Report ALCOMFT-TR-01-125, ALCOM-FT, Utrecht, May 2001. APPEARANCE: *INFORMS Journal on Computing*, 14(1):37–51, 2002.

[7] Marjan van den Akker, Han Hoogeveen, and Gerhard J. Woeginger. The two-machine open shop problem: to fit, or not to fit, that was the question. Technical Report ALCOMFT-TR-02-82, ALCOM-FT, Utrecht, May 2002. APPEARANCE: *Operations Research Letters*, 31:219–224, 2003.

[8] E. Alba, F. Almeida, M. Blesa, J. Cabeza, C. Cotta, M. Díaz, I. Dorta, J. Gabarró, C. León, J. Luna, L. Moreno, C. Pablos, J. Petit, A. Rojas, and F. Xhafa. MALLBA: A library of skeletons for combinatorial optimisation. Technical Report ALCOMFT-TR-02-120, ALCOM-FT, Barcelona, May 2002. APPEARANCE: In *Euro-Par 2002*, number 2400, pages 927–932. Springer-Verlag, LNCS, 2002.

[9] Jochen Alber, Hans L. Bodlaender, Henning Fernau, and Rolf Niedermeier. Fixed Parameter Algorithms for PLANAR DOMINATING SET and Related Problems. Technical Report ALCOMFT-TR-01-112, ALCOM-FT, Utrecht, May 2001. APPEARANCE: *Algorithmica*, 33:461–493, 2002.

[10] Susanne Albers, Lene M. Favrholdt, and Oliver Giel. On Paging with Locality of Reference. Technical Report ALCOMFT-TR-02-56, ALCOM-FT, Aarhus, May 2002. APPEARANCE: In *34th Annual ACM Symposium on the Theory of Computing*, pages 258–267, 2002.

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

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

[13] Stephen Alstrup, Gerth Stølting Brodal, and Theis Rauhe. Optimal Static Range Reporting in One Dimension. Technical Report ALCOMFT-TR-01-53, ALCOM-FT, Aarhus, May 2001. APPEARANCE: In *Proc. 33rd Annual ACM Symposium on Theory of Computing*, pages 476–482, 2001.

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

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

[16] 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.

[17] Ernst Althaus and Christian Fink. A Polyhedral Approach to Surface Reconstruction from Planar Contours. Technical Report ALCOMFT-TR-02-88, ALCOM-FT, MPI, May 2002. APPEARANCE: In *Proceedings of the 9th International Conference on Integer Programming and Combinatorial Optimization*, Cambridge, USA, May 2002.

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

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

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

[21] Maria Andreou, Dimitris Fotakis, Sotiris Nikoletseas, Vicky Papadopoulou, and Paul Spirakis. On Radiocoloring Hierarchically Specified Planar Graphs: PSPACE-completeness and Approximations. Technical Report ALCOMFT-TR-02-68, ALCOM-FT, Patras and MPI, May 2002. APPEARANCE: In *27th International Symposium on Mathematical Foundations of Computer Science, MFCS 2002*. Springer-Verlag, 2002.

[22] 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.

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

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

[25] 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.

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

[27] Lars Arge, Ulrich Meyer, Laura Toma, and Norbert Zeh. On External-Memory Planar Depth First Search. Technical Report ALCOMFT-TR-01-93, ALCOM-FT, MPI, May 2001. APPEARANCE: *Journal of Graph Algorithms and Applications*, 7(2):105–129, 2003.

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

[29] Albert Atserias. Improved Bounds for the Weak Pigeonhole Principle and Infinitely Many Primes from Weaker Axioms. Technical Report ALCOMFT-TR-02-122, ALCOM-FT, Barcelona, May 2002. APPEARANCE: *Theoretical Computer Science*, 295(1–3):27–39, 2002.

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

[31] Albert Atserias. Unsatisfiable Random Formulas are Hard to Certify. Technical Report ALCOMFT-TR-02-129, ALCOM-FT, Barcelona, May 2002. APPEARANCE: *Journal of the ACM*, 2003. To appear.

[32] 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.

[33] Albert Atserias, Maria Luisa Bonet, and Juan Luis Esteban. Lower Bounds for the Weak Pigeonhole Principle and Random Formulas Beyond Resolution. Technical Report ALCOMFT-TR-02-23, ALCOM-FT, Barcelona, May 2002. APPEARANCE: *Information and Computation*, 176(2):136–152, 2002.

[34] 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.

[35] Albert Atserias, Nicola Galesi, and Pavel Pudlák. Monotone Simulations of Nonmonotone Proofs. Technical Report ALCOMFT-TR-02-131, ALCOM-FT, Barcelona, May 2002. Appearance: *Journal of Computer and System Sciences*, 65:626–638, 2002.

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

[37] 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.

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

[39] Giorgio Ausiello, Paolo G. Franciosa, and Daniele Frigioni. Partially Dynamic Maintenance of Minimum Weight Hyperpaths. Technical Report ALCOMFT-TR-02-11, ALCOM-FT, Rome, April 2002. Appearance: *Journal of Discrete Algorithms*. to appear.

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

[41] 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.

[42] Yossi Azar, Joan Boyar, Leah Epstein, Lene M. Favrholdt, Kim S. Larsen, and Morten N. Nielsen. Fair versus Unrestricted Bin Packing. Technical Report ALCOMFT-TR-01-50, ALCOM-FT, Aarhus, May 2001. Appearance: *Algorithmica*, 34(2):181–196, 2002.

[43] Eric Bach, Joan Boyar, Leah Epstein, Lene M. Favrholdt, Tao Jiang, Kim S. Larsen, Guo-Hui Lin, and Rob van Stee. Tight Bounds on the Competitive Ratio on Accommodating Sequences. Technical Report ALCOMFT-TR-01-49, ALCOM-FT, Aarhus, May 2001. Appearance: *Journal of Scheduling*, 6(2):131–147, 2003.

[44] David A. Bader, Bernard M.E. Moret, and Peter Sanders. High-Performance Algorithm Engineering for Parallel Computation. Technical Report ALCOMFT-TR-02-29, ALCOM-FT, MPI, May 2002. Appearance: In *Experimental Algorithmics — From Algorithm Design to Robust and Efficient Software*, volume 2547 of *LNCS*, pages 1–23. Springer, 2002.

[45] R. Baeza-Yates, J. Gabarró, and X. Messeguer. Fringe Analysis of Synchronized Parallel Insertion Algorithms in 2–3 Trees. Technical Report ALCOMFT-TR-02-119, ALCOM-FT, Barcelona, May 2002. Appearance: *Theoretical Computer Science*, (299):231–271, 2003.

[46] J Baixeries and G Casas-Garriga. Sampling Strategies for Finding Frequent Sets. Technical Report ALCOMFT-TR-02-6, ALCOM-FT, Barcelona, April 2002. Appearance: In *Journées francophones d'Extraction et Gestion des Connaissances 2003*, pages 159 – 170, Lyon (France), 2003.

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

[48] 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.

[49] 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.

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

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

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

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

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

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

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

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

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

[59] Cyril Banderier, Kurt Mehlhorn, and Rene Beier. Average case analysis of three combinatorial algorithms under limited randomness. Technical Report ALCOMFT-TR-02-21, ALCOM-FT, INRIA and MPI, May 2002. Appearance: In *Proc. 28th International Symposium on Mathematical Foundations of Computer Science (MFCS-03)*, volume 2747, pages 198–207, Bratislava, Slovak Republic, 2003.

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

[61] 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.

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

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

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

[65] Hannah Bast. Scheduling at Twilight the Easy Way. Technical Report ALCOMFT-TR-02-36, ALCOM-FT, MPI, May 2002. APPEARANCE: In Afonso Ferreira and Helmut Alt, editors, *STACS 2002 : 19th Annual Symposium on Theoretical Aspects of Computer Science*, volume 2285 of *Lecture Notes in Computer Science*, pages 166–178, Antibes, Juan-Les-Pins, France, March 2002. Springer.

[66] 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.

[67] 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.

[68] 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.

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

[70] 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.

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

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

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

[74] Luca Becchetti, Stefano Leonardi, Alberto Marchetti-Spaccamela, and Kirk Pruhs. Semi-clairvoyant 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.

[75] 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.

[76] 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)*.

[77] 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.

[78] René Beier, Peter Sanders, and Naveen Sivadasan. Energy Optimal Routing in Radio Networks Using Geometric Data Structures. Technical Report ALCOMFT-TR-02-32, ALCOM-FT, MPI, May 2002. APPEARANCE: In *29th International Colloquium on Automata, Languages and Programming*, number 2380 in LNCS, pages 366–376. Springer, 2002.

[79] 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.

[80] 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.

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

[82] 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.

[83] 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.

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

[85] Petra Berenbrink, Tom Friedetzky, and Leslie Ann Goldberg. The Natural Work-Stealing Algorithm is Stable. Technical Report ALCOMFT-TR-01-72, ALCOM-FT, Warwick, May 2001. APPEARANCE: *SIAM Journal of Computing*. To appear.

[86] 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.

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

[88] 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.

[89] 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.

[90] M.J. Blesa, Ll. Hernàndez, and F. Xhafa. Parallel Skeletons for Tabu Search Method. Technical Report ALCOMFT-TR-01-152, ALCOM-FT, Barcelona, June 2001. APPEARANCE: In *8th International Conference on Parallel and Distributed Systems*, pages 23–28, Kyongju City, Korea, 2001. IEEE Computer Society Press.

[91] M.J. Blesa, Ll. Hernandez, and F. Xhafa. Parallel Skeletons for Tabu Search Method Based on Search Strategies and Neighborhood Partition. Technical Report ALCOMFT-TR-01-155, ALCOM-FT, Barcelona, June 2001. APPEARANCE: In M. Paprzycki J. Dongarra and J. Wasniewski, editors, *4th International Conference on Parallel Processing and Applied Mathematics*, volume 2328 of *Lecture Notes in Computer Science*, pages 185–193, Naleczow, Poland, September 2002. Springer–Verlag.

[92] M.J. Blesa, P. Moscato, and F. Xhafa. A Memetic Algorithm for the Minimum Weighted k-Cardinality Tree Subgraph Problem". Technical Report ALCOMFT-TR-01-165, ALCOM-FT, Barcelona, June 2001. APPEARANCE: In *4th Metaheuristics International Conference*, volume 1, pages 85–90, Porto, Portugal, 2001.

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

[94] 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.

[95] Hans L. Bodlaender. A generic NP-hardness proof for a variant of Graph Coloring. Technical Report ALCOMFT-TR-01-111, ALCOM-FT, Utrecht, May 2001. APPEARANCE: *Journal of Universal Computer Science*, 7(12):1114–1124, 2001.

[96] Hans L. Bodlaender. Necessary edges in k-chordalizations of graphs. Technical Report ALCOMFT-TR-01-113, ALCOM-FT, Utrecht, May 2001. APPEARANCE: *Journal of Combinatorial Optimization*, 7:283–290, 2003.

[97] 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.

[98] Hans L. Bodlaender, Michael J. Dinneen, and Bakhadyr Khoussainov. Relaxed Update and Partition Network Games. Technical Report ALCOMFT-TR-01-121, ALCOM-FT, Utrecht, May 2001. APPEARANCE: *Fundamenta Informaticae*, 49:301–312, 2002.

[99] 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.

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

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

[102] 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.

[103] 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.

[104] Hans L. Bodlaender and Udi Rotics. Computing the treewidth and the minimum fill-in with the modular decomposition. Technical Report ALCOMFT-TR-02-62, ALCOM-FT, Utrecht, May 2002. Appearance: *Algorithmica*, 36:375–408, 2003.

[105] Hans L. Bodlaender, Richard B. Tan, and Jan van Leeuwen. Finding a Δ-regular Supergraph of Minimum Order. Technical Report ALCOMFT-TR-01-129, ALCOM-FT, Utrecht, May 2001. Appearance: *Discrete Applied Mathematics*, 131(1):3–9, 2003.

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

[107] 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.

[108] 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.

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

[110] 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.

[111] 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.

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

[113] 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.

[114] 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.

[115] 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.

[116] Joan Boyar, Lene M. Favrholdt, Kim S. Larsen, and Morten N. Nielsen. Extending the Accommodating Function. Technical Report ALCOMFT-TR-02-4, ALCOM-FT, Aarhus, March 2002. Appearance: *Acta Informatica*, 40(1):3–35, 2003.

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

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

[119] 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.

[120] 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.

[121] 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.

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

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

[124] 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.

[125] 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.

[126] 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.

[127] 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.

[128] 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.

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

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

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

[132] 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.

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

[134] Gerth Stølting Brodal and Rolf Fagerberg. Funnel Heap - A Cache Oblivious Priority Queue. Technical Report ALCOMFT-TR-02-136, ALCOM-FT, Aarhus, June 2002. Appearance: In *Proc. 13th Annual International Symposium on Algorithms and Computation (ISAAC)*, volume 2518 of *Lecture Notes in Computer Science*, pages 219–228. Springer, 2002.

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

[136] Gerth Stølting Brodal, Rolf Fagerberg, and Christian N. S. Pedersen. Computing the Quartet Distance Between Evolutionary Trees in Time $O(n \log^2 n)$. Technical Report ALCOMFT-TR-01-131, ALCOM-FT, Aarhus, May 2001. APPEARANCE: In *Proc. 12th Annual International Symposium on Algorithms and Computation (ISAAC)*, volume 2223 of *Lecture Notes in Computer Science*, pages 731–742. Springer, 2001.

[137] Gerth Stølting Brodal, Rolf Fagerberg, and Christian N. S. Pedersen. Computing the Quartet Distance between Evolutionary Trees in Time $O(n \log n)$. Technical Report ALCOMFT-TR-02-54, ALCOM-FT, Aarhus, May 2002. APPEARANCE: *Algorithmica (Special issue on ISAAC 2001)*, 2003. To appear.

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

[139] Gerth Stølting Brodal and Riko Jacob. Time-dependent networks as models to achieve fast exact time-table queries. Technical Report ALCOMFT-TR-01-176, ALCOM-FT, Aarhus, September 2001. APPEARANCE: In *Proc. Algorithmic Methods and Models for Optimization of Railways (ATMOS 2003)*, volume 92 of *Electronic Notes in Theoretical Computer Science*, page 12. 2003.

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

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

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

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

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

[145] 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.

[146] Stefan Burkhardt and Juha Kärkkäinen. Better Filtering with Gapped q-Grams. Technical Report ALCOMFT-TR-01-74, ALCOM-FT, MPI, May 2001. APPEARANCE: *Fundamenta Informaticae*, 56(1–2):51–70, 2003.

[147] 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.

[148] 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.

[149] 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.

[150] 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.

[151] 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.

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

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

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

[155] 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.

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

[157] 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.

[158] 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.

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

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

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

[162] 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.

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

[164] 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.

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

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

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

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

[169] 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.

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

[171] 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.

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

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

[174] 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.

[175] 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.

[176] 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.

[177] 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.

[178] 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.

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

[180] 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.

[181] Ioannis Chatzigiannakis, Sotiris Nikoletseas, Nearchos Paspallis, Paul Spirakis, and Christos Zaroliagis. An Experimental Study of Basic Communication Protocols in Ad-Hoc Mobile Networks. Technical Report ALCOMFT-TR-01-119, ALCOM-FT, Patras, May 2001. Appearance: In *Algorithm Engineering – WAE 2001*, pages 159–171. Lect. Notes in Comp. Science Vol.2141 (Springer-Verlag), 2001.

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

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

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

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

[186] Ioannis Chatzigiannakis, Sotiris Nikoletseas, and Paul Spirakis. Smart Dust Local Detection and Propagation Protocols. Technical Report ALCOMFT-TR-02-114, ALCOM-FT, Patras, May 2002. APPEARANCE: In *Proceedings of the 2nd International Workshop on Principles of Mobile Computing – POMC 2002*. ACM, 2002.

[187] 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.

[188] 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.

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

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

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

[192] 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.

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

[194] 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.

[195] 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.

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

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

[198] M. Cryan, M. Dyer, L. A. Goldberg, M. Jerrum, and R. Martin. Rapidly Mixing Markov Chains for Sampling Contingency Tables with a Constant Number of Rows. Technical Report ALCOMFT-TR-02-19, ALCOM-FT, Warwick, May 2002. APPEARANCE: In *Proceedings of IEEE FOCS*, pages 711–720, 2002.

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

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

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

[202] 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.

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

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

[205] Artur Czumaj and Christian Sohler. Testing Hypergraph Coloring. Technical Report ALCOMFT-TR-01-133, ALCOM-FT, Paderborn, May 2001. APPEARANCE: In *Proceedings of the 28th International Colloquium on Automata, Languages and Programming (ICALP)*, volume 2076 of *LNCS*, pages 493–505. Springer, 2001.

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

[207] 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.

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

[209] 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.

[210] 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* 11*th European Symposium on Algorithms (ESA)*, pages 161–171, 2003.

[211] 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.

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

[213] 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.

[214] Jean-Francois Dantzer and Philippe Robert. Analysis of a multi-class queueing system. Technical Report ALCOMFT-TR-01-10, ALCOM-FT, INRIA, January 2001. APPEARANCE: *Annals of Applied Probability*, 12(1):–, 2002.

[215] 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.

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

[217] K.M.J. De Bontridder, B.V. Halldórsson, M.M. Halldórsson, J.K. Lenstra, R. Ravi, and l. Stougie. Approximation algorithms for the minimum test set problem. Technical Report ALCOMFT-TR-02-80, ALCOM-FT, Utrecht, May 2002. APPEARANCE: *Math. Programming B.*, 98:477–491, 2003.

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

[219] 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.

[220] 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.

[221] 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.

[222] 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.

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

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

[225] 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.

[226] 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.

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

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

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

[230] C. Demetrescu, I. Finocchi, G.F. Italiano, and S. Naeher. Visualization in Algorithm Engineering: Tools and Techniques. Technical Report ALCOMFT-TR-01-149, ALCOM-FT, Rome, June 2001. APPEARANCE: R. Fleischer, B. Moret, E. Meineche Schmidt (Eds.), LNCS 2547, Springer Verlag, 2002.

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

[232] C. Demetrescu, I. Finocchi, and J. T. Stasko. Specifying Algorithm Visualizations: Interesting Events or State Mapping? Technical Report ALCOMFT-TR-01-177, ALCOM-FT, Rome, October 2001. APPEARANCE: LNCS 2269, Springer Verlag, 2002.

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

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

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

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

[237] C. Demetrescu and G.F. Italiano. A New Approach to Dynamic All Pairs Shortest Paths. Technical Report ALCOMFT-TR-02-92, ALCOM-FT, Rome, May 2002. APPEARANCE: In *Proc. 35th Annual ACM Symposium on Theory of Computing (STOC'03)*, 2003.

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

[239] 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.

[240] 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.

[241] 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.

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

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

[244] 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.

[245] J. Díaz, N. Do, M. J. Serna, and N. C. Wormald. Bisection of Random Cubic Graphs. Technical Report ALCOMFT-TR-02-102, ALCOM-FT, Barcelona, May 2002. APPEARANCE: In *Randomization and approximatio techniques in Computer Science, RANDOM 2002*, Lecture Notes in Computer Science. Springer, 2002. To appear.

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

[247] J. Díaz, J. Ne ˇ setřil, M. Serna, and D.M. Thilikos. H-colorings of Large Degree Graphs. Technical Report ALCOMFT-TR-02-117, ALCOM-FT, Barcelona, May 2002. APPEARANCE: In *1st EurAsian Conference on Advances in Information and Communication Technology (EURASIA-ICT 2002)*, volume 2510 of *Lecture Notes in Computer Science*, pages 850–857. Springer-Verlag, 2002.

[248] J. Díaz, J. Petit, and M. Serna. A survey on graph layout problems. Technical Report ALCOMFT-TR-01-136, ALCOM-FT, Barcelona, June 2001. APPEARANCE: *ACM Computing Surveys*, 34(3):313–356, 2002.

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

[250] J. Díaz, M. Serna, and D. M. Thilikos. Counting H-colorings of Partial k-trees. Technical Report ALCOMFT-TR-01-142, ALCOM-FT, Barcelona, June 2001. APPEARANCE: *Theoretical Computer Science*, (281):291–309, 2002.

[251] J. Díaz, M. Serna, and D. M. Thilikos. The Complexity of Parameterized H-Colorings: A survey. Technical Report ALCOMFT-TR-01-143, ALCOM-FT, Barcelona, June 2001. APPEARANCE: In *Graphs, morphisms and statistical physics*, DIMACS series in Discrete Mathematics and Computer Science. AMS. To appear.

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

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

[254] Josep Díaz, Maria Serna, and Dimitrios Thilikos. The complexity of restrictive H-coloring. Technical Report ALCOMFT-TR-02-101, ALCOM-FT, Barcelona, May 2002. APPEARANCE: *Discrete Applied Mathematics*, 2003. to appear.

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

[256] 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.

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

[258] 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.

[259] 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.

[260] Amalia Duch and Conrado Martínez. On the Average Performance of Orthogonal Range Search in Multidimensional Data Structures. Technical Report ALCOMFT-TR-01-185, ALCOM-FT, Barcelona, November 2001. APPEARANCE: In P. Widmayer, F. Triguero, R. Morales, M. Hennessy, S. Eidenbenz, and R. Conejo, editors, *Proc. 29th Int. Col. Automata, Languages and Programming (ICALP)*, volume 2380 of *Lect. Notes in Comp. Sci.*, pages 514–524. Springer-Verlag, 2002.

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

[262] 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.

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

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

[265] Vincent Dumas, Fabrice Guillemin, and Philippe Robert. Limit results for Markovian models of TCP. Technical Report ALCOMFT-TR-01-33, ALCOM-FT, INRIA, April 2001. APPEARANCE: In *Globecom'2001*, San Antonio, November 2001.

[266] 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.

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

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

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

[270] M. Dyer, L. A. Goldberg, and M. Jerrum. Counting and sampling H-colourings. Technical Report ALCOMFT-TR-02-20, ALCOM-FT, Warwick, May 2002. APPEARANCE: *Information and Computation*. To appear.

[271] Martin Dyer, Leslie Ann Goldberg, Catherine Greenhill, Gabriel Istrate, and Mark Jerrum. Convergence of the Iterated Prisoner's Dilemma Game. Technical Report ALCOMFT-TR-01-5, ALCOM-FT, Warwick, January 2001. Appearance: *Combinatorics, Probability and Computing*, 11:135–157, 2002.

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

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

[274] 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.

[275] 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.

[276] 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.

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

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

[279] 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.

[280] Robert Elsässer, Thomas Lücking, and Burkhard Monien. New Spectral Bounds on k-Partitioning of Graphs. Technical Report ALCOMFT-TR-01-89, ALCOM-FT, Paderborn, May 2001. Appearance: In *Proceedings of the 13th ACM Symposium on Parallel Algorithms and Architectures (SPAA)*, pages 255–262. ACM, 2001.

[281] 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.

[282] 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.

[283] 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.

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

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

[286] Leah Epstein and Lene M. Favrholdt. On-Line Maximizing the Number of Items Packed in Variable-Sized Bins. Technical Report ALCOMFT-TR-02-59, ALCOM-FT, Aarhus, May 2002. APPEARANCE: *Acta Cybernetica*, 16:57–66, 2003.

[287] Leah Epstein and Lene M. Favrholdt. Optimal Non-Preemptive Semi-Online Scheduling on Two Related Machines. Technical Report ALCOMFT-TR-02-58, ALCOM-FT, Aarhus, May 2002. APPEARANCE: In *27th International Symposium on Mathematical Foundations of Computer Science*, pages 245–256, 2002.

[288] Leah Epstein and Lene M. Favrholdt. Optimal Preemptive Semi-Online Scheduling to Minimize Makespan on Two Related Machines. Technical Report ALCOMFT-TR-02-57, ALCOM-FT, Aarhus, May 2002. APPEARANCE: *Operations Research Letters*, 30(4):269–275, 2002.

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

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

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

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

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

[294] 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.

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

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

[297] 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.

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

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

[300] Lene Monrad Favrholdt and Morten Nyhave Nielsen. On-Line Edge-Coloring with a Fixed Number of Colors. Technical Report ALCOMFT-TR-01-164, ALCOM-FT, Aarhus, June 2001. APPEARANCE: In *FSTTCS: Foundations of Software Technology and Theoretical Computer Science*, volume 1974 of *Lecture Notes in Computer Science*, pages 106–116, 2000.

[301] 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.

[302] 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.

[303] 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.

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

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

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

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

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

[309] 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.

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

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

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

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

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

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

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

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

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

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

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

[321] 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.

[322] 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.

[323] 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.

[324] 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.

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

[326] 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.

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

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

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

[330] Dimitris Fotakis, Sotiris Nikoletseas, Vicky Papadopoulou, and Paul Spirakis. Radiocolorings in Periodic Planar Graphs: PSPACE-Completeness and Efficient Approximations for the Optimal Range of Frequencies. Technical Report ALCOMFT-TR-01-107, ALCOM-FT, Patras, May 2001. Appearance: In *28th International Workshop on Graph-Theoretic Concepts in Computer Science (WG 2002)*. Springer-Verlag, 2002.

[331] 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.

[332] Jens S. Frederiksen and Kim S. Larsen. Packet Bundling. Technical Report ALCOMFT-TR-02-60, ALCOM-FT, Aarhus, May 2002. Appearance: *Journal of Algorithms*, 48(2):407–428, 2003.

[333] 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.

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

[335] Christine Fricker, Philippe Robert, and Danielle Tibi. A degenerate central limit theorem for single resource loss systems. Technical Report ALCOMFT-TR-02-9, ALCOM-FT, INRIA, April 2002. APPEARANCE: *Annals of Applied Probability*, 13(2):561–575, 2003.

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

[337] Daniele Frigioni, Tobias Miller, Umberto Nanni, and Christos Zaroliagis. An Experimental Study of Dynamic Algorithms for Transitive Closure. Technical Report ALCOMFT-TR-01-102, ALCOM-FT, Patras and Rome, May 2001. APPEARANCE: *ACM Journal of Experimental Algorithmics*, 6, 2001.

[338] 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.

[339] 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.

[340] J. Gabarró, A. Stewart, and M. Clint. Grab and Go Systems: a CPO approach to concurrent web and grid-based computation. Technical Report ALCOMFT-TR-02-118, ALCOM-FT, Barcelona, May 2002. APPEARANCE: *Electronic Notes in Theoretical Computer Science*, 66(3):1–21, 2002.

[341] 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.

[342] 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.

[343] 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.

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

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

[346] 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.

[347] 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.

[348] 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.

[349] 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.

[350] 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.

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

[352] Leslie Ann Goldberg. Computation in permutation groups: counting and randomly sampling orbits. Technical Report ALCOMFT-TR-01-4, ALCOM-FT, Warwick, January 2001. APPEARANCE: In J. Hirschfeld, editor, *Surveys in Combinatorics*, pages 109–143. Cambridge University Press, 2001.

[353] Leslie Ann Goldberg, Steven Kelk, and Mike Paterson. The complexity of choosing an H-colouring (nearly) uniformly at random. Technical Report ALCOMFT-TR-02-18, ALCOM-FT, Warwick, May 2002. APPEARANCE: In *Symposium on the Theory of Computing*, pages 53–62, 2002.

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

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

[356] 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.

[357] Marco Gori and Klaus Meer. A step towards a complexity theory for dynamical systems. Technical Report ALCOMFT-TR-01-191, ALCOM-FT, Aarhus, December 2001. APPEARANCE: *Mathematical Logic Quarterly*, 48, Suppl. 1:45–58, 2002.

[358] 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.

[359] 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.

[360] Fabrice Guillemin, Philippe Robert, and Bert Zwart. AIMD algorithms and exponential functionals. Technical Report ALCOMFT-TR-02-10, ALCOM-FT, INRIA, April 2002. Appearance: *Annals of Applied Probability*, 2003. Accepted for publication.

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

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

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

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

[365] 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.

[366] 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.

[367] 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.

[368] 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.

[369] Han Hoogeveen and Gerhard J. Woeginger. Some comments on sequencing with controllable processing times. Technical Report ALCOMFT-TR-01-106, ALCOM-FT, Utrecht, May 2001. Appearance: *Computing*, 68:181–192, 2002.

[370] Sixiang Hou and Han Hoogeveen. The three-machine proportionate flow shop problem with unequal machine speeds. Technical Report ALCOMFT-TR-01-122, ALCOM-FT, Utrecht, May 2001. Appearance: *Operations Research Letters*, 31:225–231, 2003.

[371] David A. Hutchinson, Peter Sanders, and Jeffrey S. Vitter. Duality between prefetching and queued writing with Parallel Disks. Technical Report ALCOMFT-TR-01-79, ALCOM-FT, MPI, May 2001. Appearance: In *9th European Symposium on Algorithms (ESA)*, LNCS 2161, pages 62–73. Springer, 2001.

[372] 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.

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

[374] Lars Jacobsen and Kim S. Larsen. Variants of (a,b)-Trees with Relaxed Balance. Technical Report ALCOMFT-TR-01-52, ALCOM-FT, Aarhus, May 2001. APPEARANCE: *International Journal of Foundations of Computer Science*, 12(4):455–478, 2001.

[375] Lars Jacobsen, Kim S. Larsen, and Morten N. Nielsen. On the Existence and Construction of Non-Extreme (a,b)-Trees. Technical Report ALCOMFT-TR-01-78, ALCOM-FT, Aarhus, May 2001. APPEARANCE: *Information Processing Letters*, 84(2):69–73, 2002.

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

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

[378] 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.

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

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

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

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

[383] 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.

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

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

[386] 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.

[387] 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.

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

[389] Jan Klein, Jens Krokowski, Matthias Fischer, Rolf Wanka, and Friedhelm Meyer auf der Heide. The Randomized Sample Tree: A Data Structure for Externally Stored 3D-Scenes. Technical Report ALCOMFT-TR-02-64, ALCOM-FT, Paderborn, May 2002. APPEARANCE: In *Proceedings of the ACM Symposium on Virtual Reality Software and Technology (VRST '99)*, pages 137 – 146, 2002.

[390] 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.

[391] 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.

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

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

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

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

[396] Dimitrios Koukopoulos, Marios Mavronicolas, Sotiris Nikoletseas, and Paul Spirakis. On the Stability of Compositions of Universally Stable, Greedy Contention-Resolution Protocols. Technical Report ALCOMFT-TR-02-74, ALCOM-FT, Patras and Cyprus, May 2002. APPEARANCE: In *Proceedings of the 16th International Symposium on DIStributed Computing – DISC2002*, pages 88–102. Springer-Verlag, 2002.

[397] 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.

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

[399] 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.

[400] 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.

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

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

[403] 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.

[404] Christof Krick, Harald Räcke, and Matthias Westermann. Approximation Algorithms for Data Management in Networks. Technical Report ALCOMFT-TR-01-71, ALCOM-FT, Paderborn, May 2001. APPEARANCE: In *Proceedings of the 13th ACM Symposium on Parallel Algorithms and Architectures (SPAA)*, pages 237–246, 2001.

[405] 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.

[406] 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.

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

[408] 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.

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

[410] Kim S. Larsen. Relaxed Multi-Way Trees with Group Updates. Technical Report ALCOMFT-TR-01-51, ALCOM-FT, Aarhus, May 2001. APPEARANCE: *Journal of Computer and System Sciences*, 66(4):657–670, 2003.

[411] Kim S. Larsen. Relaxed Red-Black Trees with Group Updates. Technical Report ALCOMFT-TR-01-186, ALCOM-FT, Aarhus, November 2001. APPEARANCE: *Acta Informatica*, 38(8):565–586, 2002.

[412] 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.

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

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

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

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

[417] 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.

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

[419] 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.

[420] 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.

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

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

[423] 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.

[424] 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.

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

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

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

[428] J.A. Makowsky and K. Meer. Polynomials of bounded tree-width. Technical Report ALCOMFT-TR-01-36, ALCOM-FT, Aarhus, April 2001. APPEARANCE: pages 211–250, 2002.

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

[430] 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.

[431] 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.

[432] Conrado Martínez, Daniel Panario, and Alfredo Viola. Analysis of Quickfind with small subfiles. Technical Report ALCOMFT-TR-01-187, ALCOM-FT, Barcelona, November 2001. APPEARANCE: In B. Chauvin, Ph. Flajolet, D. Gardy, and A. Mokkadem, editors, *Proc. 2nd Col. on Mathematics and Computer Science (MathInfo 2002)*, Trends in Mathematics, pages 329–340. Birkhäuser Verlag, 2002.

[433] 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.

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

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

[436] K. Meer and G.W. Weber. Some aspects of studying an optimization or decision problem in different computational models. Technical Report ALCOMFT-TR-01-37, ALCOM-FT, Aarhus, April 2001. APPEARANCE: *European Journal on Operational Research*, 143, Issue 2:406–418, 2002.

[437] 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.

[438] 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.

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

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

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

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

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

[444] Ulrich Meyer. Directed Single–Source Shortest–Paths in Linear Average–Case Time. Technical Report ALCOMFT-TR-01-98, ALCOM-FT, MPI, May 2001. APPEARANCE: *Journal of Algorithms*, 48(1):91–134, 2003.

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

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

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

[448] Ulrich Meyer and Peter Sanders. Parallel Shortest Path for Arbitrary Graphs. Technical Report ALCOMFT-TR-01-95, ALCOM-FT, MPI, May 2001. APPEARANCE: *Journal of Algorithms*, 49:114–152, 2003.

[449] 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.

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

[451] 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.

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

[453] 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.

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

[455] Burkhard Monien and Robert Preis. Upper Bounds on the Bisection Width of 3- and 4-regular Graphs. Technical Report ALCOMFT-TR-01-116, ALCOM-FT, Paderborn, May 2001. APPEARANCE: In *Proceedings of 26th International Symposium on Mathematical Foundations of Computer Science (MFCS)*, volume 2136 of *LNCS*, pages 524–536. Springer, 2001.

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

[457] Jesper Makholm Nielsen. On the Number of Maximal Independent Sets in a Graph. Technical Report ALCOMFT-TR-02-87, ALCOM-FT, Aarhus, May 2002. APPEARANCE: In *Proc. 14th Symp. Discrete Algorithms*, pages 456–457. SIAM, January 2003.

[458] 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.

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

[460] 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.

[461] Sotiris Nikoletseas, Grigorios Prasinos, Paul Spirakis, and Christos Zaroliagis. Attack Propagation in Networks. Technical Report ALCOMFT-TR-01-101, ALCOM-FT, Patras, May 2001. APPEARANCE: In *Proc. 13th ACM Symposium on Parallel Algorithms and Architectures (SPAA 2001)*, pages 67–76. ACM Press, 2001.

[462] 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.

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

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

[465] 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.

[466] 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.

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

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

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

[470] Rasmus Pagh and Jakob Pagter. Optimal Time-Space Trade-Offs for Non-Comparison-Based Sorting. Technical Report ALCOMFT-TR-01-141, ALCOM-FT, Aarhus, June 2001. APPEARANCE: In *13th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 9–18. ACM Press, 2002.

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

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

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

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

[475] Jordi Petit. Layout Problems. Technical Report ALCOMFT-TR-01-157, ALCOM-FT, Barcelona, June 2001. APPEARANCE: PhD thesis, Universitat Politècnica de Catalunya, 2001.

[476] 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.

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

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

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

[480] Tobias Polzin and Siavash Vahdati Daneshmand. Extending Reduction Techniques for the Steiner Tree Problem. Technical Report ALCOMFT-TR-02-135, ALCOM-FT, MPI, June 2002. APPEARANCE: In *8th European Symposium on Algorithms (ESA)*, Lecture Notes in Computer Science, 2002. to appear.

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

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

[483] E. Meineche Schmidt R. Fleischer, B. Moret, editor. *Experimental Algorithmics : From Algorithm Design to Robust and Efficient Software*, volume 2547 of *LNCS*. Springer, 2002.

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

[485] 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.

[486] 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.

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

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

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

[490] 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.

[491] 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.

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

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

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

[495] Peter Sanders. Presenting Data from Experiments in Algorithmics. Technical Report ALCOMFT-TR-01-81, ALCOM-FT, MPI, May 2001. APPEARANCE: In *Experimental Algorithmics — From Algorithm Design to Robust and Efficient Software*, volume 2547 of *LNCS*, pages 181–196. Springer, 2002.

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

[497] Peter Sanders and Kurt Mehlhorn. Scanning Multiple Sequences Via Cache Memory. Technical Report ALCOMFT-TR-02-30, ALCOM-FT, MPI, May 2002. APPEARANCE: *Algorithmica*, 35(1):75–93, 2003.

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

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

[500] Peter Sanders and Berthold Vöcking. Random Arc Allocation and Applications. Technical Report ALCOMFT-TR-02-31, ALCOM-FT, MPI, May 2002. APPEARANCE: In *8th Scandinavian Workshop on Algorithm Theory*, number 2368 in LNCS, pages 121–130. Springer, 2001.

[501] 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.

[502] 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.

[503] 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.

[504] 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.

[505] 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.

[506] Frank Schulz, Dorothea Wagner, and Christos Zaroliagis. Using Multi-Level Graphs for Timetable Information in Railway Systems. Technical Report ALCOMFT-TR-02-70, ALCOM-FT, Patras, May 2002. APPEARANCE: In *Algorithm Engineering and Experiments – ALENEX 2002*, volume 2409, pages 43–59. Lect. Notes in Comp. Sc., 2002.

[507] Meinolf Sellmann and Torsten Fahle. Coupling Variable Fixing Algorithms for the Automatic Recording Problem. Technical Report ALCOMFT-TR-01-87, ALCOM-FT, Paderborn, May 2001. APPEARANCE: In F. Meyer auf der Heide, editor, *Proceedings of the 9th Annual European Symposium on Algorithms (ESA)*, volume 2161 of *LNCS*, pages 134–145. Springer, 2001.

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

[509] 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.

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

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

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

[513] Norbert Sensen. Lower Bounds and Exact Algorithms for the Graph Partitioning Problem using Multicommodity Flows. Technical Report ALCOMFT-TR-01-88, ALCOM-FT, Paderborn, May 2001. APPEARANCE: In F. Meyer auf der Heide, editor, *Proceedings of the 9th Annual European Symposium on Algorithms (ESA)*, volume 2161 of *LNCS*, pages 391–403. Springer, 2001.

[514] Jop Sibeyn, James Abello, and Ulrich Meyer. Heuristics for Semi-External Depth First Search on Directed Graphs. Technical Report ALCOMFT-TR-02-49, ALCOM-FT, MPI,

May 2002. Appearance: In *Proc. Symposium on Parallel Algorithms and Architectures (SPAA)*. ACM, 2002. To appear.

[515] 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.

[516] A. Stewart, M. Clint, and J. Gabarró. Algebraic rules for reasoning about BSP programs. Technical Report ALCOMFT-TR-01-156, ALCOM-FT, Barcelona, June 2001. Appearance: In S. Gorlatch and C. Lengauer, editors, *Constructive Methods for Parallel Programming*, volume 10 of *Advances in Computation: Theory and Practice*, pages 41–57. Nova Science Publishers, Inc, 2002.

[517] A. Stewart, M. Clint, J. Gabarró, and M. Serna. Towards refining BSP programs into loosely coupled distributed systems. Technical Report ALCOMFT-TR-01-139, ALCOM-FT, Barcelona, June 2001. Appearance: In R. Sakellariou, J. Keane, J. Gurd, and L. Freeman, editors, *Euro-Par 2001 Parallel Processing*, volume 2150 of *Lecture Notes in Computer Science*, pages 549–559. Springer, 2001.

[518] 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.

[519] Anna Östlin and Rasmus Pagh. One-Probe Search. Technical Report ALCOMFT-TR-02-66, ALCOM-FT, Aarhus, May 2002. Appearance: In *29th Int. Colloquium on Automata, Languages and Programming*, volume 2380 of *Lecture Notes in Computer Science*, pages 439–450. Springer, 2002.

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

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

[522] Dimitrios M. Thilikos, Maria J. Serna, and Hans L. Bodlaender. A polynomial time algorithm for the cutwidth of bounded degree graphs with small treewidth. Technical Report ALCOMFT-TR-01-115, ALCOM-FT, Barcelona and Utrecht, May 2001. Appearance: In F. Meyer auf der Heide, editor, *Proceedings 9th Annual European Symposium on Algorithms ESA 2001, Lecture Notes on Computer Science, vol. 2161*, pages 380–390. Springer-Verlag, 2001.

[523] 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.

[524] 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.

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

[526] Alexandre Tiskin. A new way to divide and conquer. Technical Report ALCOMFT-TR-01-3, ALCOM-FT, Warwick, January 2001. APPEARANCE: *Parallel Processing Letters*, 11(4):409–422, 2001.

[527] Alexandre Tiskin. Parallel convex hull computation by generalised regular sampling. Technical Report ALCOMFT-TR-01-14, ALCOM-FT, Warwick, February 2001. APPEARANCE: In B. Monien and R. Feldmann, editors, *Proceedings of Euro-Par '02*, volume 2400 of *Lecture Notes in Computer Science*, pages 392–399. Springer-Verlag, 2002.

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

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

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

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

[532] 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.

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

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

[535] 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.

[536] Michael Wand, Matthias Fischer, Ingmar Peter, Friedhelm Meyer auf der Heide, and Wolfgang Straßer. The Randomized z-Buffer Algorithm: Interactive Rendering of Highly Complex Scenes. Technical Report ALCOMFT-TR-01-134, ALCOM-FT, Paderborn, May 2001. APPEARANCE: In *Proceedings of the ACM Computer Graphics (SIGGRAPH)*, pages 361–370, 2001.

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

[538] 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.

[539] Christos Zaroliagis. Implementations and Experimental Studies of Dynamic Graph Algorithms. Technical Report ALCOMFT-TR-02-69, ALCOM-FT, Patras, May 2002. APPEARANCE: In *Experimental Algorithmics – From Algorithm Design to Robust and Efficient Software*, chapter 11, pages 229–278. Springer-Verlag, 2002.

[540] 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.