

SPECTRUM

SPECTRUM MANAGEMENT IN STATIC AND DYNAMIC NETWORKS

Department of Informatics, University of Bergen

1 Introduction

Even though existing network resources are growing at an increasing rate, the efficient use and management of these resources is paramount in being able to accommodate the even larger increase in use. This increase is twofold, both the number of users and also the volume of data each user is transmitting is increasing at an exponential rate.

In addition, new technology and services give rise to new demands that must be addressed. It follows that existing solutions already in use can become insufficient either because the size of the system has become too large or because new technology opens for new solutions not previously applicable. Thus the challenge ahead is both in studying completely new scenarios and in migrating existing solutions to new networks.

To focus and advance Norwegian research on these issues the Norwegian Research Council has allocated funds for research through the IKT-2010 program. This program has three focused areas of interest of which two are communication technologies and distributed systems. It should be noted that these two areas are not necessarily disjoint and that in fact several of the challenges in managing communication networks require distributed solutions.

Among the specific research challenges listed within the program in the area of communication technology we mention the following:

- Optimizing and making use of available frequency resources.
- Designing solutions with a large extent of flexibility and with large capacity.
- Being able to handle mobile users, terminals, and functionality.
- Being able to manage large volumes of data.

As an example of how the issues related to distributed computing overlap with communication networks we also mention the following two specific research challenges listed in the area of distributed systems: “designing adaptive software with the possibility of reconfiguration” and “designing scalable software (to millions of users)”.

Many important issues in managing advanced communication networks are inherently combinatorial and algorithmic by nature. One of the more general such challenges is as mentioned in the IKT-2010 program the management of a limited number of frequencies. This involves issues such as how to achieve the highest possible throughput, how to achieve a fair distribution of the available resources, and how to maintain quality of service in changing environments. The technologies of each system, together with national and international agreements, translate into a constrained set of usable channels to be assigned. The general procedure of effective management of the available bandwidth and frequencies is known as *spectrum management*.

Depending on the area of application some of the problems require a best possible solution due to the financial implications while others require real-time solutions that are assured to work. Most of the optimization problems related to these types of problems are known to be NP-hard and efficient optimal solutions are thus not likely to exist. Consequently, as the networks grow in size and complexity, the

need to find algorithms with better performance guarantees or run time bounds becomes increasingly more important. In addition to classical algorithms simple and efficient heuristics can also prove useful in practice.

Ad-hoc networks are perceived as one of the main technologies for future mobile networks. They have a constantly changing user group that themselves define the structure of the network. This puts extra pressure on the speed and efficiency of the underlying communication protocols. Due to their size and fluctuating structure these networks must be able to operate without any central computing unit having complete knowledge of the entire system. Still, the networks must be able to route messages and function correctly. This requires that the communication nodes execute protocols that are based only on local information. Self stabilizing algorithms are well suitable for the mentioned problems in that they operate in a distributed fashion and can accommodate a changing network structure. In addition they provide a sound mathematical foundation that complements well with a more simulation oriented approach.

The main goals of the IKT-2010 program are threefold:

1. To build fundamental competence for the future development of the commercial sector and the society in general.
2. To achieve a level of expertise among the Norwegian research groups comparable with the best international ones.
3. To transfer competence from the basic research areas to users in applied research, industry, and society in general.

The applicant group, consisting of researchers from the Department of Informatics, University of Bergen together with an industrial reference group from the main Norwegian telecom company Telenor is well qualified to participate in achieving these goals. The group consists of 5 full-time professors from the University of Bergen along with two senior scientists from Telenor. The group has a strong research focus, a well established track record of publishing internationally, and extensive international collaboration and contacts.

As we have specific competence in various aspects related to frequency assignment and distributed algorithms we propose a focused research effort on spectrum management in static and dynamic networks. In the case when one computational unit has complete knowledge of the structure of the network this includes analysis, development of algorithms and heuristics, and testing and validation on practical instances. These are research areas in which we have an extensive track record.

When there is no centralized knowledge of the network structure we will use self-stabilizing algorithms both as a tool for developing algorithms and also for testing these on real world scenarios. The group has so far produced some results in this area mainly by combining graph-algorithmic knowledge with self-stabilizing algorithms. We believe that this is a fruitful approach that will provide novel and efficient solutions for problems arising from dynamic and ad-hoc networks.

Through the close collaboration between the academic group and the industrial reference group, we will not only assert that the target problems are relevant, but also achieve an immediate feedback of the produced results to industry. In addition to the direct contact with industry, we will also be transferring knowledge by publishing our results in conferences/journals, and by supplying master and ph.d. candidates with relevant competence for the ICT-industry and civilian authorities.

2 Goals

The overall goals of the SPECTRUM project are:

G1. Advancing the state-of-the-art within spectrum management in networks

As the most important problems concerning frequency assignment in networks are NP-hard, the best algorithms currently in use are either heuristics or approximation algorithms. With the constantly

increasing size and complexity of communication networks, it is becoming even more important both to design new algorithms with better bounds on quality and time, and also to ensure that these work well in practice.

Our goal is to contribute in the form of new methods and results in the area of spectrum management. Initially we will identify the types of assignment problems, networks, and parameters of these that are of particular interest. Such parameters could be the structure and rate of change of the network and the cost, throughput, and fairness of a required solution. This process will be done in close collaboration with the industrial reference group to ensure relevance.

Based on the input from this process we will develop and apply new algorithms and heuristics for the various scenarios and evaluate how these perform over time. Since we expect several of the relevant problems to be computationally hard, it is an important decision whether to strive for an optimal solution or to apply heuristic search techniques. Thus part of the work will also be to determine for which settings there exist efficient solutions and if this is not the case to investigate if the pay-off is higher by either applying heuristics or by using various forms of intelligent brute force techniques.

We will combine self-stabilizing algorithms with solid algorithm theory to produce new results on problems connected to dynamic and ad-hoc networks. The type of problems we will be investigating are similar to those of static networks, but when one only has local information this becomes particularly challenging. Several self-stabilizing algorithms on general network topologies have an exponential worst case running time, thus it is our aim to see how one can develop more efficient algorithms both on general and more restricted network topologies. With our background in algorithms (sequential, parallel, and distributed) and optimization we believe that we are well qualified to achieve this.

G2. Strengthening the Norwegian academic community around these issues

The algorithms group in Bergen is the most internationally active algorithms research group in Norway within the mentioned areas. With the combined expertise of the applicant group and the new knowledge and results that will be gained through the SPECTRUM project, our goal is to increase the national competence around these issues.

G3. Disseminating and exchanging the knowledge acquired

The acquired knowledge is to be presented at appropriate national and international journals, conferences, and workshops.

G4. Transfer of existing and acquired knowledge to industry

By having an industrial reference group we will ensure that the particular problems investigated are of relevance and that produced results are quickly transferred back to industry. We will also be producing M.Sc. and Ph.D. graduates to be employed by the industry with a solid background in relevant issues within the topics of the SPECTRUM project. These students will not only transfer knowledge back to the industry, but will also interact and receive guidance during the project from our industrial partner.

3 Results

The following results are planned:

R1. Methods to aid in development of efficient and practical solutions of spectrum management for a network with a given structure (goal G1)

Our contribution here will be in the form of:

- Specification of new problems that model practical challenges and needs that arise due to new technology and applications.
- Applying, testing, and adapting known solution methods to new problems.

- Designing new algorithms with better bounds on the quality of the produced result, or on the run time of the algorithm.
- Designing fast search heuristics and testing these in practice.
- Proposing methods to identify special scenarios on which some of the generally intractable assignment problems can be solved optimally and efficiently.
- Establishing the trade offs between search heuristics and algorithms that finds optimal (or close to optimal) solutions.

R2. Self stabilizing algorithms for spectrum management on dynamic and mobile ad-hoc networks (goal G1)

We will combine our solid background in algorithms and complexity, and in the area of parallel and distributed computing, with the notion of self-stabilizing algorithms to design efficient algorithms for issues related to spectrum management in dynamic and ad-hoc networks. These results will rely on the following:

- Modeling and analyzing spectrum management issues related to mobile ad-hoc networks in the setting of self-stabilizing algorithms.
- Defining new methods for efficiently and effectively use of information stored both in the immediate neighborhood and also in the close vicinity of a node in a mobile ad-hoc network.
- Verification of the proposed self-stabilizing algorithms through tests on practical instances.

R3. Scientific publications in the area (goal G3)

The results achieved through the project will be published through:

- papers in international journals and international refereed conferences such as ICDCS, WG, ESA, SWAT, DISC, PODC, SODA, and OPODIS
- papers and presentations at national conferences and meetings
- presentation via the SPECTRUM project web site.

We expect to produce close to 40 international journal and conference publications within the project topics during the project period.

R4. Winter schools (goals G2, G3, and G4)

We will organize:

- a yearly winter school where Norwegian students and students from abroad get training in problem solving methods within the topics of the SPECTRUM project.

R5. Training experts in network algorithms (goals G4)

During the project period we expect to graduate 12 M.Sc. and 2 Ph.D. students in this area. We will have a seminar on related issues every semester where basic and new results are presented. In addition we will arrange problem solving sessions, where the students are trained to solve related problems and design appropriate algorithms.

R6. Bringing together industry and academia (goals G1, G2, and G4)

The industrial reference group along with the ongoing IKT-2010 project BEATS (Professor Kjell Jørgen Hole) both express the need for solid algorithmic results to aid in the challenges they meet in practical implementation of new technologies. We will cooperate closely with them to bridge the gap between hardware solutions and algorithmic software solutions. In addition, the M.Sc. and Ph.D. candidates who graduate from the SPECTRUM project are expected to join industrial companies and use and disseminate their knowledge and skills.

4 State-of-the-art, Challenges, and Needs

Frequencies in the radio spectrum represent a scarcely available resource in modern communications networks. Best possible capacity utilization, and thereby profitability, of a network is therefore closely tied with optimal spectrum management. This has inspired a substantial amount of research work on graph algorithms and combinatorial optimization models designed for various spectrum management problems in networks.

International research concerning communications networks can in many ways be viewed as two-fold. The industry develops hardware technology that gives new technological possibilities and challenges, whereas academia concentrates on how the new technology can be exploited better with new theoretical methods. Existing methods currently in use can become either insufficient or inefficient because of the new complex structure and large size enabled by the new technology, or they do not fully exploit the new possibilities presented by the new technology.

However, the fruitful cooperation between industry and academia is often hindered by a “cultural gap” between the more theoretical oriented academic researchers and engineers working in an applied industrial setting. There are several reasons for this such as different background and education, and having to meet different types of demands. As a consequence engineers working with hardware often do not have the required background and training to bring state-of-the-art academic developments into practical technological implementations. On the other hand academic researchers might simplify and idealize their models to such an extent that produced results might not always be applicable or practical in a real world setting. It is only through sustained contact and interaction between practitioners from both sides that one can assure synergy, that the studied problems are relevant, and that the produced solutions fit into a realistic framework.

To be able to handle large datasets and user groups in both static and mobile networks require that the limited spectrum resources be managed in a best possible way. Although there are many different issues involved in achieving this, at the heart of any such method there must be an algorithm or heuristic that is both effective *and* efficient. It is a common misconception that as computers become faster the importance of the complexity of the underlying algorithms becomes less significant. In fact, the opposite is true: Consider the size of the largest problem one can solve in a given amount of time. When the speed of the computer is increased the difference between what an asymptotically slow algorithm and what a fast algorithm can solve will *increase*.

The field of network and protocol design depends heavily on simulations for development and analysis of new results. However, these simulations might not always be sufficient to guarantee correct operation and even if they do, they might not scale as well when the size of network grows [51]. Thus there is a need both for verifying existing algorithms/protocols and also for developing new algorithms based on sound mathematical methods.

International academic research within the field of network algorithms is broad. Several large EU projects (ALCOM-FT, ARACNE, APPOL, COST 273) concentrate on contributing to the solution of these important problems. As many of the relevant problems have an inherently combinatorial and algorithmic nature, we find several of the international collaborators of the applicant group among the participants in these EU projects. To mention but one such collaboration, Professor Fedor Fomin’s previous position (before coming to Bergen in November 2002) was at the Heinz Nixdorf Institute, University of Paderborn, Germany, in the group headed by Professor Friedhelm Meyer auf der Heide, principal investigator of the ALCOM-FT project, see <http://www.brics.dk/ALCOM-FT/>. These studies typically involve various aspects of selected fundamental optimization issues in modern network communications, ranging from the design and theoretical analysis to experimental validation of both efficient, robust, and stable algorithmic solutions. The SPECTRUM project shares these objectives.

The frequency assignment problem occurs in a variety of networks, for example mobile phone systems, and is typically abstracted as a binary relation and modeled as a graph. This then results in a graph coloring problem, with vertices representing transmitters/receivers having edges between nearby pairs, and where a frequency assignment is obtained by mapping the colors to the linearly ordered set of frequencies [63, 64]. The standard graph coloring problem can be viewed as asking for a partition of the vertices into

a few independent sets. However, in many cases the standard graph coloring model is too limited as there are several other issues that must also be taken into account to obtain a more realistic model. Examples of such issues include:

- **co-channel interference** The power of a signal decreases approximately as the square of the distance from the source. Any signal received on a particular frequency will interfere with the use of this frequency. Thus if the signal to interference ratio on a particular frequency is too low then this frequency should not be used.
- **inter-channel interference** Due to the nature of radio signals, signals transmitted on frequencies that lie close together in the spectrum will interfere with each other. Thus transmitters that are physically close should avoid using channels that are close together in the spectrum.
- **bounded number of frequencies** The number of available frequencies in a particular application is typically bounded [24]. Thus it is not given that all requests for bandwidth at any given time can be satisfied.
- **minimum order or minimum span frequency assignment** Minimizing the number of different frequencies (MO-FAP) [41], or alternatively the bandwidth (MS-FAP) [63], used on a set of connections, such that actual interference constraints are satisfied.
- **simultaneous frequency and power assignment** In some applications not only the frequency but also the power can be selected for each connection [20]. This introduces more flexibility when taking the interference into account, and the challenge is to exploit this flexibility in an optimal way.
- **dynamic channel selection** Over time a transmitter might dynamically switch which frequency to use due to co-channel interference [55].

Complicating matters even further, it may be necessary to consider the fact that in some cases transmitters can share a common frequency by e.g. time multiplexing the channel.

In addition to model extensions imposed by the physical system, there is also the issue of which parameters are of importance. As an example an application such as real-time video requires that a certain amount of bandwidth be allocated for the duration of an entire session, while a setting where one is trying to maximize the total amount of transmitted data through a network would lead to other requirements.

Unfortunately, most problems thus defined are intractable, and one cannot expect that they can be solved efficiently (in polynomial-time) by traditional algorithm techniques. The on-line and dynamic versions of the same problems are at least as challenging as the standard versions. Therefore, much of the research in this vein must be focused on heuristics and search based methods.

Networks with a decentralized structure such as ad-hoc networks encounter many of the same challenges and problems as centralized networks but now with the added difficulty that there is no central computing facility. Thus algorithms must be distributed and based mainly on local information.

We see self-stabilizing algorithms as a way of developing, verifying, and testing algorithms for these types of networks. Self-stabilizing algorithms are based on sound mathematical ground and should complement well with a more simulation oriented approach. The algorithms operate in a distributed fashion with each node in the network only basing its actions on local information [22, 58]. One of the main advantages is that if a fault occurs the algorithm will correct itself and stabilize to a new valid state. In the same way self-stabilizing algorithms can also handle changing operating conditions.

Work on self-stabilizing algorithms for problems related to frequency assignment has so far resulted in methods for producing valid distance-1 colorings for special graph classes such as trees [68], bipartite graphs [70], planar graphs [38], and general graphs [39, 48]. To our knowledge these are the only settings in which problems related to frequency assignment have been studied in the context of self-stabilizing algorithms.

5 Contributions from the SPECTRUM Project

5.1 Focus of the project

Although there have been significant advances in obtaining higher throughput of communication through wireless channels new technical advances and requirements constantly give rise to further challenges that must be addressed. We believe that the current applicant group has the possibility to make significant contributions toward this end by developing new methods in the area of spectrum management.

First of all the composition of the group ensures that we can bridge the gap between industry and academic research. The academic group spans the field of algorithms, optimization, and distributed and parallel computing with practitioners well used to confirmation and validation through testing. Similarly the industrial reference group has close contact with the latest developments and trends in the telecom industry.

Our efforts will be focused both on settings where there is a centralized or decentralized organization of the network. The types of problems studied for both settings are the same but the approach differs. Within the first area, we will continue research along the lines of our previous expertise and contribute to advancing the state-of-the-art on relevant issues through new algorithms and heuristics. Within the second area our goal is to incorporate the notion of self-stabilizing algorithms with solid algorithmic knowledge in the design of useful algorithms for issues related to spectrum management in dynamic and ad-hoc networks. Throughout the entire process we will maintain close contact with the industrial reference group to ensure that both the addressed problems and the solutions are of practical interest.

Finally, it is of vital importance that any new suggested solutions and algorithms are experimentally validated. Assigning such implementations as Master's projects is a well-established tradition in the applicant group, and recently we have also been able to publish these experimental results at the emerging conferences on algorithm engineering [10, 37, 40, 50]. In the following we give a more detailed description of the types of methods we will be studying first for centralized and then for the decentralized setting.

Centralized networks: In this setting it is assumed that one computational unit has complete knowledge of the structure of the network and the changes that occur within it.

As many of the most important and relevant problems in spectrum management are NP-hard the existing methods have to relax the requirements of optimal solutions for all input instances. Still there is a tool chest of methods that can be applied once the model and the problem specification are clear. These include heuristics and various search methods [67], approximation algorithms [5], and algorithms that take better advantage of the structural properties inherent in the communication networks at hand [74]. As an example of the last point we note that several problems that are intractable on general graphs have polynomial time solutions on graphs such as planar graphs or radial graphs, both which are relevant for modeling wireless communication networks. We also mention the newly emerging field of parameterized algorithms that seems to be one fruitful direction for designing exact algorithms for certain hard problems [23, 26].

Decentralized networks: As mentioned in Section 4 there are often several non-local issues and also the restriction on the number of available frequencies that must be taken into account in realistic frequency assignment settings. This makes it particularly challenging to design self-stabilizing algorithms for these problems. Non-local issues requires that decisions made at any node are transmitted to the appropriate neighborhood. A restriction on the number of frequencies could lead to a situation where not all requests can be satisfied. In terms of self-stabilizing algorithms this means that the system cannot stabilize but must dynamically adjust over time to allow for some notion of fairness.

Designing and analyzing self-stabilizing algorithms is typically more difficult than that of their sequential counterparts. Where this turns out to be a problem we see two basic ways which this can be dealt with. It is perceivable that one can show an algorithm to converge faster through statistical analysis or through simulations. The other possibility is by restricting the algorithms to special types of overlay networks on the general network on which problems can be solved efficiently. For instance, work on tree networks has lead to several efficient polynomial time algorithms [12, 19]. However, there has been little effort to move self-stabilizing algorithms to other less restricted networks (or graphs) other than general networks.

Similarly there has been little effort to validate the time complexity of self-stabilizing algorithms through simulations [53].

It should be noted that the methodologies applicable to self-stabilizing algorithms are in many cases similar to those arising in parallel and distributed systems.

5.2 The Applicant Group

The applicant group has the following permanent faculty members from the Department of Informatics, University of Bergen:

- Professor Fedor Fomin, age 35, employed since November 2002.
- Associate Professor Dag Haugland, age 42, employed since 2000.
- Associate Professor Pinar Heggernes, age 34, employed since 2001.
- Professor Fredrik Manne, age 38, employed since 1998.
- Professor Jan Arne Telle, age 41, employed since 1998.

The industrial reference group will consist of the following employees from Telenor research and development:

- Senior scientist Geoffrey Canright
- Scientist Kenth Engø

Professor Fredrik Manne will function as the Principal Investigator (PI). His background both from applied industry as well as academia makes him well suited to lead the SPECTRUM project.

After graduating from the University of Bergen he spent two years working as a senior system consultant for the Norwegian oil company Norsk Hydro with a main responsibility for large scale applications. Following this he worked as a senior scientist at the University of Bergen Supercomputing center Parallab before joining the Department of Informatics in 1995 as an associate professor. He also held an adjunct position at the University of Oslo in 2000. In 2002 he was promoted to full professor and is also currently the vice chairman of the department.

Throughout his work he has kept close contact with industry and is today involved in several projects with various Norwegian companies. His main line of research has been focused on algorithms both sequential and parallel, with a preference for validating analytical results through practical simulations. Lately, his research focus has also expanded into distributed computing and various problems related to communication networks. He has graduated 7 M.Sc. students and one Ph.D. student and is currently the supervisor of 1 doctoral student and 8 M.Sc students.

For the remaining members we may mention that the international committee who evaluated Professor Fedor Fomin before joining the University of Bergen in November 2002. commented that “his is an impressive research output in the past six years, witnessing very novel and creative approaches” and that “his record indicates that his future research performance should have stellar qualities”. Besides his background from university, Dr Dag Haugland also has experience as consultant at the world’s leading software vendor in optimization, ILOG SA in Paris. His research in applied optimization includes works on modeling and algorithm design in areas like petroleum engineering, transportation, and signal processing. He was appointed chairman of the eighth meeting of the Nordic section of the Mathematical Programming Society (2002), and he is a program committee member of the Norwegian Signal Processing Symposium (2003). Dr Pinar Heggernes, the only female faculty-member in the Departments of Informatics, Mathematics, and Physics at the University of Bergen, is also recently employed and has already established herself as a talented and successful researcher. She has recently served as a program committee member of the Scandinavian Workshop on Algorithm Theory 2002 and is also the head of the organizing committee for The European Symposium on Algorithms to be arranged in Bergen 2004. Professor Jan Arne Telle has

consistently published internationally since starting his doctoral studies in 1988, appearing as co-author of about 40 refereed publications. He has functioned as Principal Investigator of several NFR-sponsored research projects, and as advisor for one completed PhD and also one ongoing PhD student.

All 5 members are internationally active researchers. Since 1998 we have a total of 77 papers either appearing in refereed journal publications or in conferences with LNCS-(type) proceedings.

In 2002 the Research Council of Norway presented an international evaluation of all academic research groups within the field of ICT. In this evaluation the the Optimization group at the University of Bergen to which Dr Haugland belongs was given the highest score of “Excellent”, while the Algorithms group to which the other applicants belong was given the second highest score of “Very Good”. It should be noted that this evaluation was performed before Dr Fomin joined the group.

5.3 Our position and how we can contribute

As described, the applicant group is well positioned to contribute toward meeting the goals of the IKT-2010 program. The group has a strong record of publishing internationally as well as contacts both at the national and international level with other research groups and with industry within the area communication networks. Thus we are able to provide a well-informed design and analysis of issues related to spectrum management including experimental validation, hopefully resulting in both efficient, robust, and stable results.

On the research level, the members of the group all have experience within the field of communication networks and in particular with issues related to spectrum management:

- Frequency Assignment [14, 27, 35, 37]
- Graph partitioning [13, 17, 18, 21, 28, 29, 42, 49, 52, 56, 57, 62, 66, 73]
- Self-stabilization [11, 12, 33, 34]

Except for the last item above, these are long-standing scientific interests of the members. The members of the applicant group also have the relevant experience with the main methodologies that are applicable in the SPECTRUM project:

- Heuristics, Algorithm Engineering and Experimental validation [1, 10, 25, 37, 40, 44, 45, 46, 47, 50, 54, 65]
- Parallel and Distributed Computing [8, 35, 36, 59, 60, 61, 69]
- Approximation and Fixed-Parameter algorithms [2, 13, 15, 31, 32, 42, 43]
- Algorithms for special graph classes [3, 4, 6, 7, 9, 15, 16, 30, 71, 72, 73]

The proposed project is therefore a natural continuation of the earlier research and work of the members of the applicant group. This is also substantiated by the fact that we have recently been involved in the following related projects:

- Self-stabilizing algorithms for ad-hoc networks. Grant from Fujitsu-Siemens. This is a joint project with Dr Hole from the IKT-2010 funded program BEATS. Dr Manne is the PI from the University of Bergen.
- ALCOM-FT: Algorithms and Complexity - Future Technologies. This is a joint effort between eleven of the leading groups in algorithms research in Europe, with network communication being one of 3 main research directions. Supported by the European Commission under the Future and Emerging Technologies part of the Information Society Technologies programme of the Fifth Framework, as project number IST-1999-14186.

- Efficient solutions for large-scale cellular phone frequency assignment problems. Travel grants 2000-2002 from the Aurora Programme, a France-Norway Collaboration Research Project of NFR and French Ministry of Education, Research and Technology. Dr Telle is the PI on the Norwegian side.
- Frequency management in mobile networks. Russian-Dutch Research Cooperation, NWO (Netherlands Organization for Scientific Research) grant 047.008.006.
- Algorithms for frequency assignments in mobile phone systems and generalizations of graph colorings. KVA (Royal Swedish Academy of Sciences) grant for cooperation between Sweden and the former Soviet Union. Dr Fomin is the PI on the Russian side.

It should be noted that none of these projects have included money for positions in Norway, e.g. for doctoral students. The SPECTRUM project is therefore our attempt to join our work in this direction to contribute also to the development of expertise within Norway on these issues. The group is currently involved in a bid to the European Union to establish a Network of Excellence in the area of “Structural and Algorithmic Aspects for algorithms for telecommunication (SAGA)”. This is joint effort by several leading research groups in Europe. It should be noted that if this project should get accepted it again does not contain any money for positions in Norway.

5.4 International research contacts

The strength of the applicant group is in its young age and its international outlook. Our international contacts are plentiful as is evident from a look at the co-authors in our publication lists.

The algorithms group (AG) at the University of Bergen is the only internationally active algorithms group in Norway. As such it has always felt a special responsibility for giving Norway a strong research profile within the field of algorithms. The AG has twice been the hosts for the biennial international conference “Scandinavian Workshop on Algorithm Theory” with about 80 participants from over 15 countries each time. In 2004 it will host ALGO 2004, a multi-conference event containing the main annual European algorithms conference ESA 2004 (European Symposium on Algorithms) organized jointly with APPROX 2004 (Approximation Algorithms for Combinatorial Optimization Problems) and WABI 2004 (Workshop on Algorithms in Bioinformatics) The AG consistently receives short-term international visitors, in particular from amongst our co-authors. During the last two years the group has also hosted a postdoctoral researcher on a 6-month NATO Science Fellowship Stipend (Dr Jiri Fiala, Charles University, Prague) and a 12-month visiting professor (Professor Jean Blair, West Point Academy). The AG hosted a workshop on Self-stabilizing algorithms in June 2002 with Professor Steve Hedetniemi from Clemson University visiting for one week. Professor Hedetniemi also hosted one of our doctoral students (Petter Kristiansen) for a 6-month period working on self-stabilizing algorithms.

The optimization group at the University of Bergen hosted the “Eighth Meeting of the Nordic Section of the Mathematical Programming Society” in September 2002, and in connection with this meeting Professor Martine Labbé from Université Libre de Bruxelles visited the department. In connection with our work on combinatorial optimization in transportation science, we received a visit from Professor Gilbert Laporte at Centre for Research on Transportation (CRT), Montréal. Professor Laporte’s colleague at CRT, Professor Michel Gendreau, is currently hosting our Ph.D. student Sin C. Ho working on heuristics for routing problems.

5.5 Impact on industry and society

Although the ICT-industry is currently not expanding as fast as was projected only some years ago it is still believed that this area, and in particular, telecommunication will have a continued growth well into the foreseeable future. The challenges that lie ahead are especially demanding in a country like Norway, as Norway has one of the highest uses in the world of mobile telephones per inhabitant. Combining this with a geographically distributed population and demanding terrain there is going to be substantial challenges in implementing the next generations of mobile networks.

The Bergen area contains several large companies working in the telecom area: Telenor, the leading Norwegian telecommunication company, along with its subsidiaries has large branch offices in Bergen. NextGenTel, Norway's second largest supplier of broadband services has its main office in Bergen. Nera, the leading Norwegian company in development of wireless broadband equipment is located in Bergen. In addition, together with Stavanger, the Bergen area also contains the bulk of the Norwegian oil and gas industry. Both on and offshore these rely heavily on modern ICT-solutions in their operation. Thus there are several takers for candidates from the SPECTRUM project.

The applicant group supplies candidates to the Norwegian ICT-industry both on the master and Ph.D. level. Since 1998 the group has produced 3 Ph.D. and 18 M.Sc. students in Informatics and are currently supervising 24 M.Sc. students and 3 doctoral students. Although the ICT-industry has decreased their hiring since the end of the .com area, our students have never experienced problems in attaining jobs after graduation.

Two additional doctoral students would join the group within the SPECTRUM project. As of today (summer 2003) we have several potential candidates to fill such positions.

We believe that it is of importance that candidates who leave academia to work in the industry both have a strong theoretical background and also know how to make use of this in practical implementations. These are the type of candidates we expect to produce in the SPECTRUM project.

In addition to supplying candidates to the Norwegian ICT-industry we will also achieve transfer of knowledge through our collaboration with the Research and Development group in Telenor within the area of mobile ad-hoc networks.

6 Work Packages

W1. Project management

This package covers project coordination and reporting.

(T1) **Task: Project coordination** There will be monthly meetings among the project participants for communication of the achieved results and planning.

Deliverables: Internal reports, work notes, drafts of reports and articles, presentations.

(T2) **Task: Project reporting**

Deliverables: Annual reports to NFR, status reports, updates of the project web site.

W2. Supervision package (result R5)

This package covers the supervision tasks within the project.

(T1) **Task: Supervising Ph.D. 1**

(T2) **Task: Supervising Ph.D. 2**

(T3) **Task: M.Sc. supervision**

The M.Sc. students will receive weekly supervision. For every Ph.D. student there will be one senior staff member involved in the project made responsible for the supervision. The supervision will be reviewed in the project meetings, if necessary.

W3. Analysis and modeling of problems related to spectrum management (results R1, R2, R3, R4, R6)

(T1) **Task: Defining and analyzing problems on general networks**

Research challenges: Practical problems that arise in the implementation of new technology must be properly formulated in formal models in order to be addressed by algorithmic methods.

Deliverables: A series of problem definitions and network models based on practical specifications provided through collaboration with our industrial partner Telenor.

- (T2) **Task: Identifying networks on which problems can be solved efficiently**
Research challenges: Generally intractable problems can sometimes be solved efficiently on special types of networks defined by the structural properties inherent in the communication networks at hand. In order to benefit from this, it is important to define and identify new types of networks with this property, or to design algorithms that can efficiently recognize already known special types of networks having such properties.
Deliverables: Scientific papers in conferences and journals containing proofs that special networks have desirable properties and new algorithms for recognizing such networks.
- (T3) **Task: Complexity analysis of problems on dynamic and mobile ad-hoc networks**
Research challenges: Problems are naturally harder to solve on dynamically changing networks than on static networks. In addition, problems that have efficient solutions on static networks may be intractable on dynamic networks. It is important to establish the computational complexity of problems on mobile ad-hoc networks.
Deliverables: Scientific papers in conferences and journals disseminating the advances and results obtained in this topic.

W4. Efficient heuristics and algorithms for spectrum management (results R1, R2, R3, R4, R6)

- (T1) **Task: Heuristics and search methods**
Research challenges: As most relevant problems within the field of spectrum management are NP-hard there is a constant need to develop new efficient heuristics for the various problems at hand. **Deliverables:** New heuristics and search methods as well as application of known methods to new problems. The results will be presented through scientific papers in conferences and journals.
- (T2) **Task: Approximation algorithms and practical optimal algorithms**
Research challenges: To supplement the use of heuristics one can use approximation algorithms or various clever brute force techniques to solve hard problems related to spectrum management. Thus there is a need for new approximation algorithms with better guarantees on the produced solution as well as better guarantees on runtime. In addition one must determine for what instances it is possible to devise practical optimal algorithms.
Deliverables: New approximation algorithms and optimal solution algorithms with fast runtimes in practice, presented through scientific papers in conferences and journals.
- (T3) **Task: Efficient self-stabilizing algorithms for dynamic and mobile ad-hoc networks**
Research challenges: Join our solid background in algorithm theory, and parallel and distributed computing, with the method of self-stabilizing algorithms to design efficient algorithms with provable guarantees.
Deliverables: Efficient self-stabilizing algorithms for dynamic and ad-hoc networks verified through theory and implementations. New techniques and algorithms presented through scientific papers in conferences and journals.

W5. Software packages (results R1, R2, R3, R4, R5, R6)

- (T1) **Task: Computer programs with test sets and test results**
Research challenges: Bridging the gap between theory and practice, following an idea all the way through from theoretical analysis, via algorithm design, to the final product in form of software programs. In this work package, M.Sc. students will be heavily involved.
Deliverables: Computer programs, web applets, test sets, M.Sc. candidates.

W6. Dissemination and contacts with industry and other partners (results R3, R4, R5, R6)

- (T1) **Task: Scientific presentation**
The results of the project will be disseminated through the normal scientific channels.
Deliverables: Papers in journals and presentations at national and international conferences.

(T2) **Task: Dissemination through teaching**

Teaching of students at all levels will be an important form of transferring the gained knowledge to users outside academia.

Deliverables: Two Ph.D. candidates, about 12 M.Sc. candidates.

(T3) **Task: Winter schools**

Deliverables: A winter school every year as long as the project lasts.

Work schedule

Preliminary timetable, given as project year/quarter, for the tasks and results of the project.

year/quarter	2004				2005				2006				2007			
Work package	1	2	3	4	1	2	3	4	1	2	3	4	1	2	3	4
W1	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x
W2	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x
W3	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x
W4			x	x	x	x	x	x	x	x	x	x	x	x	x	x
W5					x	x	x	x	x	x	x	x	x	x	x	x
W6	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x	x

7 Resources

The project is planned to start January 2004 and run for 4 years. The main cost of the project will be the salaries for the researchers involved: 2 Ph.D. students (3 years each) and 1 postdoctoral position for one year. In addition there is a need for funding for equipment, travel, as well as visiting researchers. The members of the applicant group, the 5 UiB professors and associate professors, will contribute with 33% of their work time each. The cost of the various parts are given in the following table (in 1000 NOK).

Overall costs (in 1000 NOK)	2004	2005	2006	2007	Sum
Doctoral student	267.5	535	535	267.5	1605
+ Equipment	50				50
+ Visit abroad		50			50
Doctoral student		535	535	535	1605
+ Equipment		50			50
+ Visit abroad			50		50
Postdocs		316.5	316.5		633
+ Equipment		50			50
Visits to Bergen	50	50	50	50	200
Conference travel	221	222	222	222	887
Equipment	30	30	30	30	120
Visits between Telenor and UiB	20	20	20	20	80
Winter school	30	30	30	30	120
Total applied for from NFR	668.5	1888.5	1788.5	1154.5	5500
33% time of 3 UiB professors (funded by UiB)	750	750	750	750	3000
33% time of 2 UiB assoc. prof. (funded by UiB)	400	400	400	400	1600
Project total	1818.5	3038.5	2938.5	2304.5	10100

7.1 New researchers

One of the main goals of the project is to educate new researchers through a doctoral program. We have therefore included 2 positions for doctoral students. To get the best possible candidates for these positions we wish the first to start no later than the summer of 2004 and the second in January 2005. All

previous doctoral students advised by the members of the applicant group have spent at least 6 months abroad visiting another research institution as part of their doctoral studies. This is something we wish to continue, and we have therefore included funding for a 6 months stay abroad for each student. We have also included funding for a postdoctoral position for one year.

For both doctoral students and the postdoc we have included a one-time expense for funding of computer equipment.

7.2 International network

To be able to produce and disseminate scientific results it is vital that we are able to maintain and nourish our international contacts. This is done both in terms of travels *from* Bergen to conferences/workshops and other research institutions and also by bringing other researchers *to* Bergen for shorter or longer stays.

We have estimated around 2 conference/workshop travels for each of the 8 participants of the project every year. In addition we have included a yearly expense for mobile computer equipment that is necessary for the project.

Both in 2002 and in 2003, the Algorithms Group at the Department of Informatics, UiB, arranged a winter school in algorithms, which has been a great success. This was attended not only by students and researchers from the University of Bergen but also by visitors from the Czech republic, Germany, and USA. We were able to do this organization thanks to extraordinary funds which were available to us only for a limited period, and we will not be able to continue this activity with our regular research resources. As this was an extremely positive and fruitful event, we would like to establish this as a part of the SPECTRUM project.

References

- [1] S. O. AASE, R. NYGAARD, J. H. HUSØY, AND D. HAUGLAND, *Optimized time and frequency domain methods for ecg signal compression*, Applied Signal Processing, 5 (1998), pp. 210–225.
- [2] B. ASPVALL, M. HALLDORSSON, AND F. MANNE, *Approximations for the generalized block distribution of a matrix*, Theoretical Computer Science, 262 (2001), pp. 145–160.
- [3] B. ASPVALL AND P. HEGGERNES, *Finding minimum height elimination trees for interval graphs in polynomial time*, BIT, 34 (1994), pp. 484–509.
- [4] B. ASPVALL, A. PROSKUROWSKI, AND J. A. TELLE, *Memory requirements for table computations in partial k -tree algorithms*, Algorithmica, 27 (2000), pp. 382–394.
- [5] G. AUSIELLO, P. CRESCENZI, G. GAMBOSI, V. KANN, AND A. M.-S. AND M. PROTASI, *Complexity and Approximation*, Springer, 1999.
- [6] A. BERRY, J.-P. BORDAT, AND P. HEGGERNES, *Recognizing weakly triangulated graphs by edge separability*, Nordic Journal of Computing, 7 (2000), pp. 164–177.
- [7] A. BERRY, P. HEGGERNES, AND Y. VILLANGER, *An on-line incremental approach for dynamically maintaining chordal graphs*, (2003). Submitted to ISAAC 2003 - International symposium on algorithms and computation.
- [8] P. BJØRSTAD, F. MANNE, T. SØREVIK, AND M. VAJTERŠIĆ, *Efficient matrix multiplication on SIMD computers*, SIAM J. Matrix Anal. Appl., 13 (1992), pp. 386–401.
- [9] J. BLAIR, P. HEGGERNES, S. HORTON, AND F. MANNE, *Broadcast domination algorithms for interval graphs, series-parallel graphs, and trees*, Tech. Rep. 249, University of Bergen, Norway, 2003. Reports in Informatics.

- [10] J. R. S. BLAIR, P. HEGGERNES, AND J. A. TELLE, *A practical algorithm for making filled graphs minimal*, Theoretical Computer Science A, 250 (2001), pp. 125–141.
- [11] J. R. S. BLAIR AND F. MANNE, *Efficient multi stage self-stabilizing algorithms for tree networks*. In preparation for journal submission.
- [12] ———, *Efficient self-stabilizing algorithms for tree networks*, in proceedings of ICDCS'03, the 23rd International Conference on Distributed Computing Systems, IEEE, Computer Society, 2003.
- [13] H. BODLAENDER, J. GUSTEDT, AND J. A. TELLE, *Linear-time register allocation for a fixed number of registers*, in Proceedings SODA'98, Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, 1998.
- [14] H. L. BODLAENDER, H. BROERSMA, F. V. FOMIN, A. V. PYATKIN, AND G. J. WOEGINGER, *Radio labeling with pre-assigned frequencies*, in Proceedings of the 10th European Symposium on Algorithms (ESA 2002), Springer, 2002.
- [15] H. L. BODLAENDER AND F. V. FOMIN, *Approximation of pathwidth of outerplanar graphs*, J. of Algorithms. to appear.
- [16] ———, *Tree decompositions with small cost*, in Proceedings of the 8th Scandinavian Workshop on Algorithm Theory (SWAT 2002), Springer, 2002.
- [17] H. BROERSMA, F. V. FOMIN, P. A. GOLOVACH, AND G. J. WOEGINGER, *Backbone colorings for networks*, in Proceedings of the 29th Workshop on Graph Theoretic Concepts in Computer Science (WG 2003), Springer, 2003.
- [18] H. BROERSMA, F. V. FOMIN, J. KRATOCHVÍL, AND G. J. WOEGINGER, *Planar graph coloring with forbidden subgraphs: Why trees and paths are dangerous*, in Proceedings of the 8th Scandinavian Workshop on Algorithm Theory (SWAT 2002), Springer, 2002.
- [19] S. C. BRUELL, S. GHOSH, M. H. KARAATA, AND S. V. PEMMARAJU, *Self-stabilizing algorithms for finding centers and medians of trees*, SIAM Journal on Computing, 29 (1999), pp. 600–614.
- [20] G. DAHL, K. JÖRNSTEN, G. LØVNES, AND S. SVAET, *Graph optimization problems in connection with the management of mobile communications systems*, Telecommunication Systems, 3 (1995), pp. 319–339.
- [21] E. D. DEMAINE, F. V. FOMIN, M. T. HAJIAGHAYI, AND D. M. THILIKOS, *Fixed-parameter algorithms for the (k,r) -center in planar graphs and map graphs*, in To appear in the proceedings of the 30th International Colloquium on Automata, Languages and Programming (ICALP 2003), Springer, 2003.
- [22] E. W. DIJKSTRA, *Self-stabilizing systems in spite of distributed control*, CACM, 17 (1974), pp. 643–644.
- [23] R. G. DOWNEY AND M. R. FELLOWS, *Parameterized Complexity*, Springer, 1999.
- [24] A. EISENBLÄTTER, M. GRÖTSCHEL, AND A. M. C. A. KOSTER, *Frequency planning and ramifications of coloring*, 2000.
- [25] O. K. ENDRESEN AND D. HAUGLAND, *Optimal adjustment of surfaces to point sets*, in Proceedings of Nordic MPS '02, Tech. Report No. 237, Department of Informatics, University of Bergen, 2002.
- [26] M. R. FELLOWS, *Parameterized complexity: The main ideas and some research frontiers*, in ISAAC, vol. 2223 of Lecture Notes in Computer Science, Springer, 2001, pp. 291–307.

- [27] J. FIALA, A. V. FISHKIN, AND F. V. FOMIN, *Online and offline distance constrained labeling of disk graphs*, in Proceedings of the 9th Annual European Symposium on Algorithms (ESA 2001), F. M. auf der Heide, ed., Springer, 2001, pp. 464–476. Lecture Notes in Computer Science, vol. 2161.
- [28] J. FIALA, P. HEGGERNES, P. KRISTIANSEN, AND J. A. TELLE, *Generalized H-coloring and H-covering of trees*, in Proceedings of 28th International Workshop on Graph-Theoretic Concepts in Computer Science, Springer, 2002.
- [29] F. FOMIN, P. HEGGERNES, AND J. A. TELLE, *Graph searching, elimination trees, and a generalization of bandwidth*, 2003. To be presented at FCT 2003 - International Symposium on Fundamentals of Computation Theory.
- [30] F. V. FOMIN, *Pathwidth of planar and line graphs*, Graphs and Combinatorics. to appear.
- [31] F. V. FOMIN, D. KRATSCHEK, AND J. C. NOVELLI, *Approximating minimum cocolorings*, in Proceedings of the 13-th International Symposium on Fundamentals of Computation Theory (FCT 2001), Springer, 2001, pp. 118–125. Lecture Notes in Computer Science, vol. 2138.
- [32] F. V. FOMIN AND A. LINGAS, *Approximation algorithms for time-dependent orienteering*, Inform. Proc. Letters. to appear.
- [33] M. GAIRING, R. M. GEIST, S. T. HEDETNIEMI, AND P. KRISTIANSEN, *A self-stabilizing algorithm for maximal 2-packing*, Tech. Rep. 230, Department of Informatics, University of Bergen, 2002. Submitted.
- [34] M. GAIRING, S. T. HEDETNIEMI, P. KRISTIANSEN, AND A. A. MCRAE, *Self-stabilizing algorithms for $\{k\}$ -domination*, in Proceedings of the Symposium on Self-stabilization, SSS-2003 (to appear), Springer, 2003.
- [35] A. GEBREMEDHIN, I. GUERRIN, J. GUSTEDT, AND J. A. TELLE, *Graph coloring on a coarse grained multiprocessor*, Discrete Applied Mathematics. to appear in special issue, editor Dieter Kratsch.
- [36] ———, *Pro: A model for parallel resource-optimal computation*, in Proceedings HPCS'2002 - The 16th Annual International Symposium on High Performance Computing Systems and Applications, 2002.
- [37] A. GEBREMEDHIN AND F. MANNE, *Scalable parallel graph coloring algorithms*, Concurrency: Practice and Experience, 12 (2000), pp. 1131–1146.
- [38] S. GHOSH AND M. H. KARAATA, *A self-stabilizing algorithm for coloring planar graphs*, Distributed Computing, 7 (1993), pp. 55–59.
- [39] M. GRADINARIU AND S. TIXEUIL, *Self-stabilizing vertex coloring of arbitrary graphs*, in Proc. OPODIS'2000, 4th international conference on principles of distributed systems, 2000, pp. 55–70.
- [40] J. GUSTEDT, O. MÆHLE, AND J. A. TELLE, *The treewidth of java programs*, in Proceedings ALENEX'02- 4th Workshop on Algorithm Engineering and Experiments, 2002.
- [41] W. K. HALE, *Frequency assignment: Theory and applications*, in Proceedings of the IEEE, vol. 68, 1980, pp. 1497–1514.
- [42] M. HALLDORSSON, G. KORSATZ, A. PROSKUROWSKI, R. SALMAN, H. SHACHNAI, AND J. A. TELLE, *Multi-coloring trees*, in Proceedings COCOON'99 - Fifth Annual International Computing and Combinatorics Conference, Springer, 1999, pp. 271–280.
- [43] M. HALLDORSSON, J. KRATOCHVIL, AND J. A. TELLE, *Independent sets with domination constraints*, Discrete Applied Mathematics, 99 (1999), pp. 39–54.

- [44] D. HAUGLAND, J. G. HEBER, AND J. H. HUSØY, *Compressing data by shortest path methods*, in Operations Research Proceedings, Springer, 1997, pp. 145–150.
- [45] ———, *Optimization algorithms for ecg data compression*, Medical and Biological Engineering and Computing, 35 (1997), pp. 420–424.
- [46] D. HAUGLAND AND S. STORØY, *A combinatorial optimization model for vector selection in frame based signal representation*, Tech. Rep. 221, Department of Informatics, University of Bergen, 2001.
- [47] ———, *Local search methods for ℓ_1 -minimization in a data compression problem*, Tech. Rep. 235, Department of Informatics, University of Bergen, 2002.
- [48] S. T. HEDETNIEMI, D. P. JACOBS, AND P. K. SRIMANI, *Self-stabilizing maximal independent set and greedy coloring in linear time*. Submitted.
- [49] P. HEGGERNES AND J. A. TELLE, *Partitioning graphs into generalized dominating sets*, Nordic Journal of Computing, 5 (1998), pp. 128–142.
- [50] P. HEGGERNES AND Y. VILLANGER, *Efficient implementation of a minimal triangulation algorithm*, in Proceedings of 10th European Symposium on Algorithms (Engineering and Applications Track), Springer, 2002.
- [51] J. HEIDEMANN, K. MILLS, AND S. KUMAR, *Expanding confidence in network simulations*, IEEE Network, (2001), pp. 58–63.
- [52] M. R. HENZINGER AND J. A. TELLE, *Faster algorithms for the nonemptiness of street automata and for communication protocol pruning*, in Proceedings SWAT'96 - Fifth Scandinavian Workshop on Algorithm Theory, Springer, 1996.
- [53] T. HERMAN, *Self-stabilization bibliography: Access guide*. Chicago Journal of Theoretical Computer Science, Working Paper WP-1, initiated November 1996. <http://www.cs.uiowa.edu/ftp/selfstab/bibliography/>.
- [54] S. HO AND D. HAUGLAND, *A tabu search heuristic for the vehicle routing problem with time windows and split deliveries*, Computers & Operations Research, (2003). In press.
- [55] I. KATZELA AND M. NAGHSHINEH, *Channel assignment schemes for cellular mobile telecommunication systems: A comprehensive survey*, IEEE Personal Communications, (1996), pp. 10–31.
- [56] J. KRATOCHVIL, A. PROSKUROWSKI, AND J. A. TELLE, *Complexity of graph covering problems*, Nordic Journal of Computing, 5 (1998), pp. 173–195.
- [57] P. KRISTIANSEN AND J. A. TELLE, *Generalized h-coloring of graphs*, in Proceedings ISAAC'00 - 7th Annual International Symposium on Algorithms And Computation, Springer, 2000.
- [58] L. LAMPORT, *Solved problems, unsolved problems, and non-problems in concurrency*, in Proceedings of the 3rd Annual ACM Symposium on Principles of Distributed Computing, 1984, pp. 1–11.
- [59] V. LO, S. RAJOPADHYE, S. GUPTA, D. KELDSSEN, M. MOHAMED, B. NITZBERG, J. A. TELLE, AND X. ZHONG, *Oregami: Tools for mapping parallel computations to parallel architectures*, International Journal of Parallel Programming, 20 (1991), pp. 237–270.
- [60] V. LO, S. RAJOPADHYE, J. A. TELLE, AND X. ZHONG, *Parallel divide-and-conquer on meshes*, IEEE Transactions on Parallel and Distributed Systems, 7 (1996), pp. 1049–1059.
- [61] F. MANNE AND H. HAFSTEINSSON, *Efficient sparse Cholesky factorization on a parallel SIMD computer*, SIAM J. Sci. Comput., 16 (1995), pp. 934–950.
- [62] F. MANNE AND T. SØREVIK, *Optimal partitioning of sequences*, J. Alg., 19 (1995), pp. 235–249.

- [63] R. A. MURPHEY, P. M. PARDALOS, AND M. G. C. RESENDE, *Frequency assignment problems*, in Handbook of combinatorial optimization, Supplement Vol. A, Kluwer Acad. Publ., Dordrecht, 1999, pp. 295–377.
- [64] L. NARAYANAN AND S. M. SHENDE, *Static frequency assignment in cellular networks*, Algorithmica, 29 (2001), pp. 396–409.
- [65] R. NYGAARD AND D. HAUGLAND, *Signal compression by piecewise linear non-interpolating approximation*, in Proceedings of International Conference on Acoustics, Speech and Signal Processing (ICASSP), Seattle, Washington, March 1999, pp. 1273–1276.
- [66] B. OLSTAD AND F. MANNE, *Efficient partitioning of sequences*, IEEE Trans. Comput., 44 (1995), pp. 1322–1326.
- [67] C. R. REEVES, *Modern Heuristic Techniques for Combinatorial Optimization*, Blackwell, Oxford, 1993.
- [68] S. K. SHUKLA, D. ROSENKRANTZ, AND S. S. RAVI, *Observations on self-stabilizing graph algorithms for anonymous networks*, in Proceedings of the Second Workshop on Self-stabilizing Systems, 1995, pp. 7.1–7.15.
- [69] T. SØREVIK, F. MANNE, R. MOE, AND A. GEBREMEDHIN, eds., *Applied Parallel Computing, Proceedings PARA 2000*, Springer, 2000. Lecture Notes in Computer Science 1947.
- [70] S. SUR AND P. K. SRIMANI, *A self-stabilizing algorithm for coloring bipartite graphs*, Information Sciences, (1992), pp. 219–227.
- [71] J. A. TELLE, *Tree-decompositions of small pathwidth*, Discrete Applied Mathematics. to appear.
- [72] J. A. TELLE AND A. PROSKUROWSKI, *Efficient sets in partial k -trees*, Discrete Applied Mathematics, 44 (1992), pp. 109–117.
- [73] ———, *Algorithms for vertex partitioning problems on partial k -trees*, SIAM Journal on Discrete Mathematics, 10 (1997), pp. 529–550.
- [74] J. VAN LEEUWEN, *Handbook of Theoretical Computer Science, volume A, chapter Graph Algorithms - Classes of graphs*, Elsevier, 1990.