MADALGO
CENTER FOR MASSIVE DATA ALGORITHMICS

Aarhus University

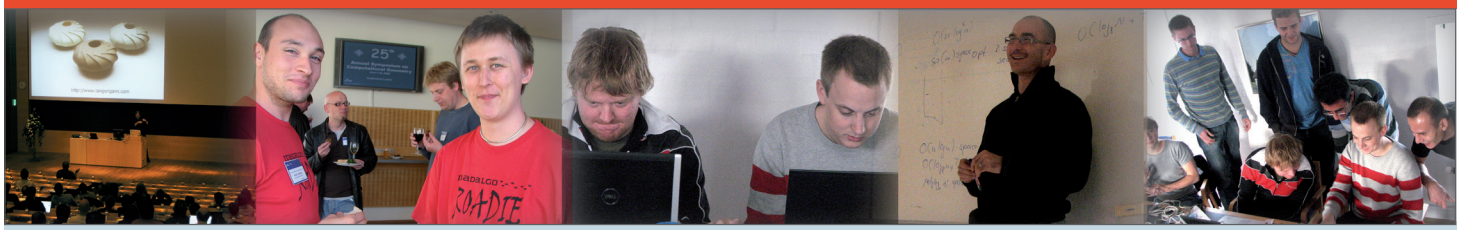

ANNUAL REPORT 2009

CENTER FOR MASSIVE DATA ALGORITHMICS

MADALGO - a Center of the Danish National Research Foundation / Department of Computer Science / Aarhus University

MADALGO
CENTER FOR MASSIVE DATA ALGORITHMICS

Aarhus University

2009 Highlights

Research team

At the end of 2009 the center research team consisted of six senior researchers (2 at AU), four Post Docs (3 at AU) and eighteen PhD students (12 at AU). Additionally, three further Post Docs and 3 PhD students (who obtained degrees during the year) were part of the center in 2009. All center Post Docs are internationals and so are a good deal of the PhD students.

Research collaboration and results

In 2009 MADALGO researchers published 60 peer reviewed research paper within the center research areas. Several of these papers have appeared in highly ranked journals and conference proceedings. Some of the results in the papers have been obtained with the many international researchers that have visited MADALGO in 2009.
The center also has extensive multidisciplinary and industry collaboration, e.g. on issues in connection with massive terrain data. In fact, in 2009 Aarhus University and several MADALGO researchers established the company SCALGO to commercialize software for efficient massive terrain data processing.

MADALGO researchers consider algorithms in many different theoretical models designed to capture different aspects of real computer hardware.

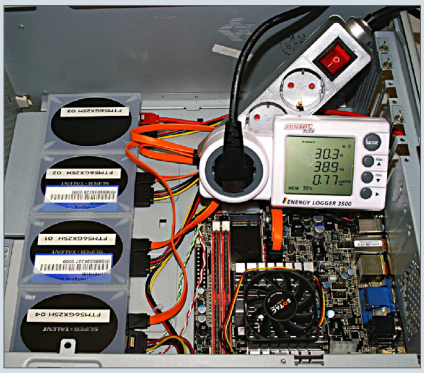

For obvious reasons, energy efficient computing is attracting increasing attention. By combining and engineering algorithm techniques from several theoretical models, MADALGO researchers recently obtained a world record in the area of energy efficient computing: Sorting 1 Terabyte of data with minimal energy use.

Center events

In 2009 MADALGO organized the 25th Annual ACM Symposium on Computational Geometry, which is the top international conference in the area. The day before the symposium MADALGO organized a celebration of the 25th anniversary of the symposium with lectures by four distinguished invited speakers.

The day after the symposium MADALGO organized the First Workshop on Massive Data Algorithmics. The workshop was a success and the goal is to eventually make it a full-fledged annual conference.

In 2009 center researchers also gave more than 20 presentations at international research conferences, participated in more than 15 invitation only international research workshops, and gave more than 25 invited presentations at research workshops and University seminars.

Awards and acknowledgments

The senior center researchers received a number of awards and acknowledgments in 2009. Arge was named a Distinguished Scientist by the Association of Computing Machinery (ACM) and appointed Associate Editor of the ACM Journal of Experimental Algorithmics. Demaine was named the Mathematical Association of America's George Polya Lecturer, and Brodal was promoted to Associate professor with exceptional qualifications ("Lektor MSK").

Content

This report describes the 2009 activities at the Danish National Research Foundation *Center for Massive Data Algorithmics* (MADALGO). The outline of the report more or less follows the agenda of the annual review meeting for the center held on March 18, 2010 (except that external funding and teaching is only covered in the appendices). The report is accompanied by a number of appendices (covering external relations, conferences, educational activities, academic services, external funding, awards, public outreach, patents and applications, publications, list of personnel) as specified by the foundation; note that *some* of these appendices only cover the employees at University of Aarhus (and not the participants at Max Planck Institute for Informatics, Massachusetts Institute of Technology and Frankfurt University). Finally, note that the 2009 accounts for the center with appendices are also important in order to obtain a complete overview of the 2009 activities of the center.

Center director statement

By signature it is confirmed that the annual report and accounts/budget with accompanying notes and appendices contain all relevant information regarding the annual primary activities in the Danish National Research Foundation Center for Massive Data Algorithmics.

31 *March, 2010* Lars Arge

Lars Arge
Center Director

1 Center background and organization

Center for Massive Data Algorithmics (MADALGO) is motivated by the spectacular increase in the amount of data being collected and processed in many modern applications, as well as the inadequacy of traditional algorithms theory in many such applications – in particular that traditional theory has used a simplistic machine model that does take the hierarchical memory organization of modern machines into account. The memory system of a typical computer system is made up of several levels of cache, main memory and disk, where the access times of the different levels can vary by several orders of magnitude, and where data is transferred between the levels in large contiguous blocks. This means that often it is much more important to minimize the number of block transfers than it is to minimize CPU computation, as is done in the simplistic models used in traditional algorithmics. Thus, the inadequacy of the simplistic machine models often translates into software inadequacies when processing massive data.

MADALGO strives to become a world-leading center in algorithms for handling massive data, where massive is interpreted broadly to cover computations where the data is large compared to the resources of the computational device. The high-level objectives of the center are:

- To significantly advance the fundamental algorithms knowledge in the area of efficient processing of massive datasets
- To train the next generation of researchers in a world-leading and international environment
- To be a catalyst for multidisciplinary collaborative research on massive dataset issues in commercial and scientific applications

To meet these objectives the center builds on the research strength at the main center site at Aarhus University (AU) in Denmark (with Brodal and Arge as core faculty), at the center sites at the Max Planck Institute for Informatics (MPI) and at Frankfurt University (FRA) in Germany (with Mehlhorn and Meyer as core faculty), and at Massachusetts Institute of Technology (MIT) in the US (with Demaine and Indyk as core faculty). The center also relies on international research collaboration, multidisciplinary and industry collaboration, and in general on a vibrant international environment at the main center site. Another key is the focus on three related but also very different research focus areas, namely *I/O-efficient*, *cache-oblivious* and *streaming algorithms*, as well as *algorithm engineering*. The four center sites complement each other nicely in terms of research strength within these core focus areas. Please refer to the original research plan for a more thorough discussion of the center motivation, objectives, core research areas and research strength at the various sites. The four focus areas will also be described further in Section 2.

Organizationally the center has been quite stable during 2009. Scientifically the center continues to be lead by center director Lars Arge along with the other core faculty and with advice from the centers international advisory board. Logistically the main center site continues to be managed by center manager Else Magård and accountant Ellen Lindstrøm (half time), with secretarial support from the Department of Computer Science. Two programmers (one founded by the center) have also been associated with the center in 2009. On the research personnel side the center developed more or less as planned in 2009. During the year an Associate Professor position was advertized at AU in connection with the center. The applicants for the position are under evaluation and it is expected that the position will be filled by August 2010. The hire should add significant strength in the core center research areas. Center Post Docs and PhD-students are discussed further in Section 5; refer also to the Appendices for a 2009 list of personnel.

2 Center research activities

2.1 Research plan

The original center research plan discusses a number of main research challenges in each of the four focus areas, and lists a number of concrete open geometric and graph problems. However, it is also noted that the list of problems is non-exhaustive, and that the outcome of current and future research of course should influence the exact directions taken in the center. Furthermore, the plan discusses that as the center matures, other interesting methodologies for massive data processing than I/O-efficient, cache-oblivious and streaming algorithms should be considered. In the 2007 and 2008 annual report we described several such new methodologies/directions. Below we quickly recall all of these areas and highlight some main problems.

Main focus areas

I/O-efficient algorithms are algorithms designed in a two-level memory-disk external memory (or I/O-) model, where the memory hierarchy consists of a main memory of limited size M and an external memory (disk) of unlimited size; the goal is to minimize the number of times a block of B consecutive elements is read (or written) from (to) disk (an I/O-operation, or simply I/O). The model is motivated by the fact that transfers between main memory and disk, rather than e.g. CPU computation, is the bottleneck when processing massive datasets residing on disk. The research plan outlines a number of open problems in the area of I/O-efficient algorithms, including a number of fundamental geometric data structure and graph traversal problems, as well as some very practically motivated terrain data processing problems.

Cache-oblivious algorithms are algorithms designed in the I/O-model – but without knowledge of M and B – and then analyzed as I/O-model algorithms. Since the I/O-model analysis holds for any block and memory size, it holds simultaneously on all levels of any multi-level memory hierarchy. Thus the cache-oblivious model is effectively a way of modeling a complicated (maybe even unknown and/or changing) multi-level hierarchy using the simple two-level I/O-model. The cache-oblivious algorithms area is relatively new, with many even very fundamental problems remaining open. The research plan outlines a number of fundamental problems to be considered, mainly on geometric data structure and batched problems.

Streaming algorithms are algorithms designed in a model where only one (or a small constant number of) sequential pass(es) over the data is (are) allowed. The goal is to solve a given problem while using significantly smaller space than the input data size, and while processing each data object as fast as possible. The model is motivated by the fact that when processing truly massive datasets, solutions requiring more than one sequential pass over the data are often infeasible, since random access to disk blocks are much slower than sequential access. Moreover, in some applications data simply has to be processed sequentially as it is generated. The research plan outlines a number of fundamental problems to be considered, including investigation of the general applicability of existing streaming algorithm design techniques.

Algorithm engineering covers the design and analysis of practical algorithms, efficient implementation of these algorithms, as well as experimentation that provide insight into their applicability and further improvements. Algorithm engineering is naturally an integrated part of the center both because a main motivation for the center is the inadequacy of traditional algorithms theory in providing practically efficient algorithms, and because engineering work naturally supports multidisciplinary and industry collaboration. Additionally, algorithm engineering work often provide valuable input to theoretical model work. The research plan outlines a number of algorithm engineering challenges in the I/O-efficient and cache-oblivious algorithms areas. In the cache-oblivious area focus is on engineering simple algorithms for very fundamental problems, whereas in the I/O-efficient algorithms area focus is on leveraging existing basic I/O-algorithms libraries to further engineer algorithms for fundamental problems and to (further) develop software for efficient processing of massive terrain data.

New directions

Chip manufactures are increasingly producing chips with several CPUs (or cores) on a single chip. Thus there is a need for parallel algorithms that can use all of these cores. In the previous year's annual reports we outlined some challenges and results in a *parallel private-cache model*, which is basically a parallel extension of the I/O-model where P processors have a main memory of limited size M each and share an external memory of unlimited size. The goal is to minimize the number of parallel I/Os, in which blocks of B elements can simultaneously be transferred between the external memory and each of the main memories.

Another hardware trend is that *flash memory* devices are becoming increasingly large and cheap. Flash memory may eventually replace disks as the external storage in e.g. mobile computing. Since flash memory have very different characteristics than both internal memory and disks, we in the previous reports described work on characterizing the new devices in order to develop models of flash memory, as well as initial work on efficient algorithms for these models/devices.

Modern memory is not always fully reliable. Sometimes the content of a memory location may be temporarily or permanently corrupted, and error rates are expected to increase as memory is getting smaller and more complex. When processing massive data one usually performs a large number of memory accesses over a long period of time. Thus memory errors can become a serious problem. In the previous reports we discussed work in a recently proposed interesting model for modeling and handling memory errors called the *faulty-memory RAM model*.

When handling massive data it is often important to use really space-efficient data structures. *Succinct data structures* are data structures that support efficient queries while occupying an amount of space that is provably close to the information-theoretic minimum. In the previous reports we also discussed work on such succinct data structures.
Finally, in last year's annual report we discussed initial work on combining the different methodologies/models described above. The goal of this "*model combination"* work is of course to develop better and more realistic models for efficient algorithms design. Another crosscutting "theme" that we described was the increasing center *emphasis on data structures*. Not only is much of the work done in the center in the various modes described above really on data structuring problems, but the broad data structure expertise in the center also represents an opportunity to consider even classical open data structure problems.

2.2 Overview of selected 2009 results

In general 2009 was a busy year with steady progress on both the original research plan problems and on problems in the new direction areas outlined above. In fact, the number of center research papers published in 2009 was higher than the number of papers published in 2007 and 2008 combined. In this section we briefly discuss some of the obtained results in each of the main focus research areas, as well as in the new directions areas; whenever relevant new directions results will be discussed in connection with main focus area results. Note that due to space limitations, we will not be able to give a complete overview of all the results obtained by center researchers; see Appendix I for a full list of papers published in 2009.

I/O-efficient algorithms

During 2009 we have obtained a number of interesting results in relation to the I/O-model problems outlined in the research plan, as well as in the updated research plans in the 2007 and 2008 yearly reports. Some of the results are not only I/O-efficient but cuts across the areas discussed in Section 2.1.

Our work on I/O-efficient geometric algorithms has focused on problems in terrain data analysis and on solutions to fundamental geometric problems in the parallel private-cache model.
In the area of terrain analysis, we have obtained results on so-called terrain flooding, which is an important step in computing realistic river networks. Traditional methods for computing river networks do not account for water collecting in the basins of the terrain and, thus, yield unrealistic flow patterns. More realistic predictions of water flow across the terrain are obtained either using ``local flooding'' approaches, which remove ``small'' basins before computing the river network, or computing *exactly* when the basins spill over based on their watersheds and volume. As also briefly mentioned in last year's report, in 2009 we have presented an I/O-efficient local flooding algorithm that can remove basins based on their height, area, or volume [C85]. This result extends our earlier result on removing basins based on their height, and provides a basis for more realistic approximations of the real flooding behavior of a given terrain. In work to be presented in the 2010 Symposium on Computational Geometry, we describe an I/O-efficient algorithm that computes the spill times of all basins of the terrain *exactly*. This algorithm requires a more global analysis of the terrain than used by local flooding techniques and is more costly than our local flooding algorithm. It remains to be investigated if the added accuracy justifies the added cost.
In the parallel private-cache model, we investigated which I/O-efficient design techniques can be used to obtain I/O-efficient parallel solutions to fundamental geometric problems. For some geometric problems, such as convex hull, the I/O-efficient solution is based on parallel techniques, and it is not surprising that I/O-efficiency and parallelism can be combined quite easily for these problems. For other problems, such as line segment intersection, the I/O-efficient solutions seem inherently sequential, while existing parallel solutions are not I/O-efficient. In a paper under submission, we show how to combine the sequential I/O-efficient solutions with techniques from our previous work on parallel private-cache model algorithms [C33] to obtain efficient algorithms for orthogonal line segment intersection. Our current research, lead by parallel private-cache model expert Sitchinava, who was hired as a center Post Doc in 2009, focuses on extending this solution to arbitrary line segments.

In the graph algorithms area we have also obtained results in the parallel private-cache model using I/O-efficient techniques, while also continuing our work on I/O-efficient graph traversal related problems. The parallel private-cache results, to appear at the 2010 International Parallel and Distributed Processing Symposium (and briefly mentioned in last year's report), consists of adaption of several I/O-algorithm techniques for solving graph problems efficiently to the parallel setting resulting in several new efficient algorithms. In general, connections between parallel and I/O-efficient graph-traversal is discussed in a recent

center survey [C63]. The main I/O-efficient graph traversal related results of 2009 is an algorithm for topologically sorting a directed acyclic graph, assuming one is given a vertex-disjoint path cover of an acyclic supergraph of the graph consisting of a number (related to the main memory and block size) of directed paths. For some classes of directed acyclic graphs we show that such a path cover can be obtained I/O-efficiently as well. This work is described in a paper under submission. A recent detailed center survey provides an overview of our achievements concerning I/O-efficient traversal of general graphs [O4].

Apart from the parallel private-cache model work mentioned above, we have also presented a couple of other papers in 2009 (and also mention briefly in the 2008 report) with results that relate (or rather combine) the I/O-model with another methodologies/models. One of the papers combines the I/O and succinct data structure areas and presents the first space- and query-optimal external memory search tree structure [C76]. Various tradeoffs between time and space are also presented. The other paper investigates the interplay between I/O-efficiency and faulty-memory, and prove trade-offs between I/O-efficiency and fault tolerance for comparison based dictionaries, priority queues, and sorting [C69]. A key contribution here is the definition of formal models for simultaneously studying I/O and memory-faults. The PhD thesis of center student Mølhave gives a nice overview of our faulty-memory and terrain data processing results [T13].

Cache-oblivious algorithms

While cache-obliviousness provides an elegant framework for obtaining algorithms for multi-level memory hierarchies while reasoning in the I/O model, the techniques to obtain cache-oblivious algorithms and data structures and the limitations of the model are still poorly understood. Our research in 2009 led to several important results in both directions.

In work that was presented at the 2010 Symposium on Discrete Algorithms we obtained an optimal cache-oblivious dictionary, that is, a data structure that allows for storing and searching a collection of elements and does so using the optimal number of cache misses for element insertions and deletions and search queries. To obtain this structure, it was necessary to depart from the by now standard approach of recursively laying out a binary search tree in memory. Instead, we used a recursively constructed hierarchy of sorted buffers that are connected to form a directed acyclic graph. In 2009 we also presented a general technique for constructing cache-oblivious range reporting data structures for a wide range of query ranges, including three-sided queries in 2-d and dominance and halfspace queries in 3-d [C81]. The main challenge in constructing such data structures is obtaining an approximation of the number of points in the query range. We solved this problem cache-obliviously using so-called shallow cuttings. This result is new even in internal memory models, where no structure achieving the optimal query bound in the worst case was known before. Our new three-sided range searching data structure, just as previous structures, uses more space than I/O-efficient data structures for the problem. This raises the question whether this is unavoidable. We managed to give an affirmative answer to this question by showing that any cache-oblivious data structure that achieves the optimal query bound for any of a variety of range reporting problems must use superlinear space [C82]. This result provides a strong separation between the I/O-model and the cache-oblivious model but still leaves a substantial gap between the lower bound we were able to prove and the space usage of existing cache-oblivious range reporting data structures. Very recently center Post Doc (and data structures expert) Afshani, who was hired during 2009, along with center 2009 sabbatical visitor Zeh, have made progress on strengthening the construction to obtain even stronger lower bounds.

Streaming algorithms

As discussed in the previous year's annual reports, almost all known streaming algorithms rely of so-called sketches. In 2009 we have continued our work on computing efficient sketches of data (notably, high-dimensional vectors x) in a way that preserves important information about the data. Most of the work has been focused on obtaining a succinct approximate representation of the data, by acquiring a small number of *nonadaptive linear* measurements of x. More formally, for any "signal" represented by a vector x of length n, the representation is equal to Ax, where A is an $m \times n$ matrix. The vector Ax is often referred to as the *sketch* or the *measurement vector* of x. Although m is typically much smaller than n, the short sketch Ax contains plenty of useful information about the signal vector x. In particular, the sketch of x retains enough inherent information that we can directly recover a *sparse approximation*, i.e., an approximation of the signal that (i) contains few non-zero coefficients but (ii) nevertheless captures "most" of the information in the signal. The linearity of the sketching process turned out to be essential not only in data stream computations, but also in several other (very diverse) applications such as e.g. compressive sensing.

The new results in 2009 include a new algorithm for efficient and accurate linear sketching [C88]. Similarly to some of our earlier algorithms it provides (a) asymptotically optimal compression guarantee and (b) algorithms with near-linear running time. However, the new algorithm performs much better in practice, yielding results comparable to the state-of-the-art convex minimization algorithm (which, however, has substantially higher running time, and therefore cannot be applied to solving large problems). In work that was presented at the 2010 Symposium on Discrete Algorithms, we have also shown that that the asymptotically optimal compression guarantee (a) is optimal even for randomized compression methods. Previous results applied only to deterministic methods, or other more specific approaches. This shows that the current state-of-the art methods cannot be substantially improved. However, in other work we have shown that improvements *can* be obtained if the signals are not fully general, but instead come from restricted (but still widely applicable) classes of signals [C56]. In particular, we show that if the "large" coefficients in the signal occur close to each other, then one can increase the quality of compression significantly. We have also shown that compression quality can be substantially improved using non-linear compression [C55]. Finally, we have investigated the use of sketching techniques for tasks other than computing sparse approximation. In particular, we have developed efficient sketches that preserve the so-called Earth Mover Distance (EMD) between two images [C89], which has been shown to be a very useful metric for pattern recognition purposes. In a paper presented at the 2010 Symposium on Discrete Algorithms, we have also proposed better (in fact, near-optimal) sketches for estimating Lp norms between two vectors.

Algorithm engineering

We have continued to make progress on our libraries for efficient implementation of I/O-efficient algorithms, on implementation of I/O-efficient graph algorithms, as well as on algorithms for massive terrain data. Unfortunately, we have yet to make real progress on cache-oblivious algorithm engineering.

Our 2009 work included extensive work on the TPIE library and presentations of results (which were discussed briefly in the 2008 annual report) on both parallel-external features of the STXXL library [C65] and on an implementation of I/O-efficient shortest path computations with bounded edge-weights [C62]. The latter result, our BFS-implementation [C64], and our previous theoretical achievements (see annual reports from 2007 and 2008) are important ingredients in ongoing work towards implementations of I/O-efficient diameter approximation methods. Other ongoing graph algorithm engineering work includes work on implementation of I/O-efficient algorithms for manipulation of so-called Binary Decision Diagrams (BDDs) used extensively in verification applications.

We have continued our work on massive terrain data processing within the TerraSTREAM software project. During the year, a lot of work has been done on bug fixes, as well on adding new smaller features many of which were the result of feedback from biology and industry collaborators. Using the software and very detailed terrain data from industry collaborator COWI A/S, Biology PhD student Moeslund also conducted a study of the impact of sea-level rise on Aarhus [C78]. The study received quite a lot of attention by the local media. Our interdisciplinary work on flood prediction using very detailed terrain models also motivated much of the more theoretical work on I/O-efficient terrain flooding algorithms described above.

In last year's annual report we mentioned that a possible commercialization of the TerraSTREAM software was being explored. As a result of the continued purchasing interest from industry collaborators and others, along with the inability of AU to provide the development and support resources needed to commercialize the software, it was decided during 2009 to establish the company SCALGO ApS (short for Scalable Algorithmics ApS) to commercialize TerraSTREAM. SCALGO was established, and is owned by, center director Arge, center PhD student Revsbæk, former center PhD student (and current Duke University Post Doc) Mølhave, Duke University Faculty Agarwal and AU (and thus indirectly the Danish National Research Foundation). Revsbæk has taken a one year leave-of-absence to lead SCALGO. The company has already performed several terrain data processing consulting projects, mostly in relation to flood risk modeling.

Engineering considerations are also naturally an integrated part of much of our other work, including in the new directions areas. In fact, during 2009 we had one particular successful engineering effort that combines our work in I/O-efficient algorithms, parallel processing and flash memory, namely energy efficient sorting. Computers (and their cooling devices) have become a major factor in the consumption of electrical energy. Also, battery lifetime is the main limiting factor for many applications of mobile devices. Hence, reducing energy consumption is now an important economical and environmental goal. Using a low power processor, solid state (flash) disks, and efficient algorithm techniques from several areas, we recently managed to beat the current energy efficient sorting world record. More precicely, we beat the current records in the JouleSort benchmark (www.sortbenchmark.org) for sorting 10GB to 1TB of data by a factor of 3 and 4, respectively.

New direction areas

Above we have already discussed some of our 2009 work on parallel private-cache, flash memory and faulty-memory algorithms, as well as on succinct data structures, that also involves one or more of the core research areas, that is, that cuts across the individual models/areas. Apart from the work on *parallel private-cache* geometric and graph algorithms described above, our parallel algorithms work has also included a study of the relationship between P-completeness (a measure of the non-existence of efficient parallel solutions to a problem) and the (non-)existence of I/O-efficient solutions, with the aim to better understand the relationship between parallel and I/O-efficient algorithms [C74]. Apart from the *flash memory* work done as part of our energy efficient sorting work, we have also continued our flash memory modeling work on various levels. During the year, our work on various theoretical models of flash memory (briefly mentioned in the 2008 report) was presented at an international conference [C61]. We have also proposed a variation of standard online paging in two-level memory systems geared towards flash memory devices [C59]. In the *faulty-memory* algorithms areas we, in addition to the already mentioned work on the trade-off between I/O-efficiency and fault tolerance, studied the trade-off between the accuracy of maintaining a collection of binary counters under memory corruptions and the increment time of a counter [C60]. We studied both randomized and deterministic schemes, and showed how a randomized scheme actually allows us to achieve a running time that is independent of the number of memory-faults that actually has happened. Finally, in addition to the space optimal external search tree described above, we have also obtained a few other results on *succinct data structures*. For example (as briefly discussed in last year's report) we have described a data structure for the so-called fully indexable dictionary problem, i.e. the problem of storing a bit-vector supporting e.g. the operation of selecting the *i*'th bit in the vector, that exhibits an interesting trade-off between time and space [C77]. We have also developed asymptotically space-optimal so-called fully-functional cardinal trees [C75].

As discussed in Section 2.1, much of the work in the center focus and new directions areas is on data structures problem; above we e.g. discussed cache-oblivious [C81, C82], fault tolerant [C60] and succinct [C75, C77] data structures. We are also considering data structure problems in more classical (internal memory) models of computation, including very classical problems such as the so-called dynamic optimality conjecture,which has been open for more than 20 years. As discussed in last year's report, we have recently given a new perspective on the problem (although without solving it) [100]. This year we have also considered a classical and longstanding open problem, namely higher-dimensional orthogonal range reporting, that is, the problems of storing a set of *d*-dimensional points such that the points in an axis-aligned query hyper-box can be reported efficiently. Previously a query and space optimal solution was only known for the two-dimensional version of the problem, and it was not know if the query time must increase with the number of dimensions. We managed to develop a space and query optimal three-dimensional structure and show that the query time indeed has to increase with dimension. The work opens up some intriguing new problems. For example, the optimal two- and three-dimensional structures answers queries in the same time, but we have proven that the same time cannot be achieved in six dimensions; the question is then if the query complexity jumps at four, five or six dimensions. The work is to be presented in the upcoming Symposium on Computational Geometry and builds on some of your previous 2009 work [80].

We have also considered various other one- and two-dimensional range searching problems. In one dimension we have presented space optimal data structures for returning the elements in a subrange of the array in sorted order in linear time [C69], data structures for returning the median element in any subarray [C66], and for finding the most frequent element (mode) in any subarray. It is worth noting that the latter result (still to be submitted) is co-authored by four center PhD students. In two dimensions we have studied the problem of storing a matrix of values such that the minimum element in a rectangular query region can be reported efficiently [C57]; some of our results on this problem still have to be submitted for publication. Another two-dimensional problem we studied is the problem of storing a set of random points in the plane such that 3-sided range reporting queries can be reported using only double-logarithmic query times [C68]; some of our results on this problem will be presented at the upcoming Conference on Database Theory.

2.3 Updated research plan

As outlined in the previous section, center researchers have obtained a number of interesting and important results on the problems outlined in the original research plan, as well as on problems in the new direction areas described in the 2007 and 2008 annual reports. In general we will continue to follow the research plan, while of course letting the outcome of current and future research, as well as the interests of center researchers, guide the exact research directions taken by the center. As mentioned in Section 1, we expect the upcoming center AU faculty hire to add significant strength in the core center research areas, and thus also

result in increasing progress on the core research objectives. We will also continue to pursue the new direction areas to varying degree, focusing less on the succinct and faulty-memory areas and more on the parallel private-cache area. This choice is partly guided by research challenges (or maybe rather progress) but also by personnel: With the departure of Post Doc Rao and graduation of PhD students Mølhave and Jørgensen, the center does not really house expertise in the two former areas anymore. On the other hand, with the hire of Post Doc Sitchinava we have added significant expertise in the parallel private-cache area. In fact, with the increasing emphasis from hardware manufacturers on adding more and more cores (CPUs) on a single chip, we believe the areas will become increasingly important in massive data computation. As the number of cores grow and the memory hierarchies and interconnectivity become increasingly complex, a theoretical algorithm foundation for multicore computation, and maybe more importantly, theoretical models that capture the essential features of multicore systems in a simply yet accurate way, will be needed. Such models and algorithms are likely to be significantly influenced by current work on I/O-efficient and cache-oblivious algorithms, or more generally, by work that cuts across the different models and methodologies considered in the center. Thus we obviously plan on continuing our work in this area.

We also plan to further intensify the focus on data structure problems. As described above, the center has added Post Doc data structure expertise in 2009, and two upcoming Post Doc hires will add further expertise. One of these new Post Docs is especially interested in fundamental data structure limitations (lower bounds) and will create a strong tie to non-center theoretical computer science faculty at AU and collaborators at Tsinghua University (where he is currently working) interested in related questions. The other new Post Doc is also interested in data structure questions that relates to database problems. This fits very well with AU having just hired a world-leading database researcher (Professor Christian S. Jensen from Aalborg University) who will build up a world-class spatial database group at AU. The center work on I/O-efficient and cache-oblivious data structures already overlap with work being done in the database community, and we plan to intensify work on database related problems in these areas, in particular geometric and moving data (temporal) data structure problems.

As previous years, we also plan to continue to consider new and interesting methodologies for massive data processing. One particular area that we are currently considering, and which extends our work on models that cuts across the individual methodologies we have considered, is sensor network computation. Finally, as discussed further in the next section, we also plan on intensifying our multidisciplinary collaboration with biology and geology researchers.

3 Collaboration

The center continues to try to maintain a vibrant, world-class and international environment at the main center site, e.g. through emphasis on hosting international visitors (faculty as well as PhD students) at AU. All core MIT, MPI and FRA faculty have visited AU during 2009, and most MIT and FRA PhD-students and Post Docs have also visited. Additionally, non-center faculty Jan van Leeuwen (Utrecht) and Norbert Zeh (Dalhousie) visited AU for a longer period of time (more than 3 weeks) in 2009, along with PhD student Martin Smerek (Masaryk). The list of shorter term non-center faculty visiting AU include Rolf Fagerberg (Southern Denmark), Herman Haverkort (TU Eindhoven), Jeff Phillips (Duke), Sariel Har-Peled (UIUC), Jack Snoeyink (UNC), Martin Isenburg (LLNL), Bardia Sadri (Toronto), Ke Yi (HKUST), Riko Jacob (Munich), Christian Knauer (FU Berlin), Ian Munro (Waterloo), Rasmus Pagh (ITU) and Morteza Monemizadeh (Dortmund). In general, center researchers have extensive international research collaboration with other computer scientists, which e.g. can be seen from the list of affiliations of 2009 center collaborators with joint papers, which include Duke, Waterloo, Dalhousie, UC Irvine, NYU, Rutgers, Toronto, McGill, Stanford, Brown, Carleton, Utah, Rice, SUNY, Michigan, Berkeley, Arizona State, IBM Almaden, AT&T, Google, Bruxelles, Madrid, FU Berlin, TU Dortmund, TU Munich, TU Eindhoven, Patras, Leicester, Pisa, Paderborn, Karlsruhe, Barcelona, Weizman, Tel Aviv, Sydney, Seoul National, JAIST and Sharif. A more complete list of ongoing collaboration can be found in Appendix A.

The center also continues to try to be a catalyst for multidisciplinary and industry collaboration. Many of the centers activities in this direction are in connection with massive terrain data, where center researchers e.g. collaborate with environmental and agricultural researchers at the Department of Biolog and the Faculty of Agricultural Sciences, as well as with industry partners COWI A/S, EIVA A/S and obviously SCALGO ApS. The activities are also supported through other sources, such as a grant from the Danish strategic research council, just as center researchers are part of a recent massive terrain data processing project (with Duke and NCSU researchers) funded by the US Army Research Office. See again Appendix A for a more comprehensive list of multidisciplinary and industry collaborators.

The center is continuously seeking to extend existing collaborative efforts, as well as initiate new ones. For example, the center has been a central player in efforts to establish a *Center for Interdisciplinary Geospatial Informatics Research* (CiGIR) at AU (possibly co-located with the center), where computer scientists, biologists, geologists and others will collaborate on geospatial modeling questions, such as those appearing in biodiversity research. AU (through the University Research Foundation AUFF) has recently provided seed-funding for the effort, which will be used to hire two post docs for a two year period. One of these Post Docs will be located at MADALGO. In general, AU is currently undergoing a restructuring process with emphasis on multidisciplinary research, and it is anticipated that this process will facilitate a further strengthening of the centers multidisciplinary collaborations within the University. Other new multidisciplinary collaborations being explored include projects in connection with the truly massive data that will be generated by the future "Facility for Antiproton and Ion Research" (FAIR) and "European Spallation Source" (ESS) experiments. As mentioned in Section 2, the center also expects increased collaboration with database researchers as a result of the establishment of a strong database group at AU, just as collaboration with other theoretical computer science faculty at AU and at Institute for Theoretical Computer Science (ITCS) at Tsinghua University in Beijing (the leading research institution in China) is expected if a proposed "center for theory of interactive computing" between the two sites become a reality. As mentioned in last year report, the possibilities for collaboration with ITCS were explored during a weeklong joint ITCS and AU workshop (www.itcs.tsinghua.edu.cn/CTACC2009/) in Beijing in March 2009.

4 Events

During 2009 the center participated in and/or organized a large number of research events. These included internal weekly seminars at AU, as well as a number of informal workshops around the TerraSTREAM software. Externally, center researchers gave more than 20 presentations at international research conferences, participated in more than 15 invitation only international research workshops, and gave more than 25 invited presentations at research conferences and University seminars.

In June 2009 the center organized the 25th Annual Symposium on Computational Geometry (SoCG) at AU (see www.madalgo.au.dk/socg2009). The three day symposium, attended by 136 researchers and featuring one invited and 44 contributed talks, is the top international conference in the area of computational geometry. The day before the symposium MADALGO organized a celebration of the 25th anniversary of the symposium with invited lectures by David Dobkin (Princeton), Micha Sharir (Tel Aviv), Emo Welzl (ETH) and Kurt Mehlhorn (MPI) on the development of the field during the last 25 years. The day after the symposium MADALGO organized the first Workshop on Massive Data Algorithmics (MASSIVE - see www.madalgo.au.dk/massive2009). The aim of this new workshop is to provide a forum for researchers from both academia and industry interested in algorithms for massive dataset problems. The symposium featured 15 contributed talks and was attended by 63 researchers. Following the success of the first MASSIVE, a second workshop is being organized in connection with the 2010 SoCG. The hope is to eventually make MASSIVE a full-fledged conference. In general, judging from the evaluation conducted at the end of SoCG and MASSIVE, the three events were very successful.
During 2009 MADALGO also sponsored ALGO 2009 at the IT-University in Copenhagen, which combines the Annual European Symposium on Algorithms (ESA) with a number of smaller, specialized conferences and workshops. ALGO is the major European algorithms event and included an invited talk by MADALGO core researcher Demaine. During the year, core researchers Arge and Indyk also gave invited talks (tutorials) at the Symposium on Principles of Database Systems and at the EPFL summer institute, respectively. Center researchers have also participated in several public outreach activities. Brodal has for example lectured on massive data algorithms and algorithms in general at several high-school related events, and Arge has lectured on flood prediction using massive terrain models at several public events.

5 Research education

As mentioned, one key goal of the center is to train the next generation of researchers in a world-leading and international environment. Thus PhD-students and Post Docs are a very important part of the center, and the center strives to have a large population of international PhD students and Post Docs at AU, just as AU center students stay 6 months abroad.
Currently, the center houses 19 PhD students (12 at AU, 3 internationals) and 4 Post Docs (3 at AU, all internationals); one of these PhD students have been hired in 2010 (and thus do not appear on the 2009 list of

personnel) and one is currently on a one year leave-of-absence (and is listed with a October 31 hiring ending period). Furthermore, three center PhD students (two at AU and one at MIT) finished their PhD studies in 2009. Below we give a brief overview of the 2009 center Post Docs and PhD students, with emphasis on the AU Post Docs and PhD students. Apart from the listed PhD students, 5 MS students have also been associated with the center in 2009.

Post Docs:

- *Srinivas Rao*, PhD Chennai 2002; employed August 2007-January 2009
  Mainly works on succinct data structures. Now an Assistant Professor at Seoul National University.
- *Mohammad Adam*, PhD Eindhoven 2007; employed January 2008-December 2009
  Mainly works in computational geometry. Now a Post Doc at University of Dortmund.
- *Peter Hachenberger,* PhD MPI 2006; employed July 2008-June 2009
  Mainly works on computational geometry algorithm engineering. Now industry job.
- *Deepak Ajwani,* PhD MPI 2008; hired October 2008
  Mainly works on I/O-efficient graph algorithms and algorithms for flash memory.
- *Payman Afshani*, PhD Waterloo 2008; hired February 2009.
  Mainly works on geometric data structures.
- Nodari Sitchinava, PhD UC Irvine 2009; hired August 2009.
  Mainly works on parallel private-cache model algorithms.

Additionally, *Gabriel Moruz* (PhD AU, 2007) continues to be a Post Doc at FRA. He mainly works on fault tolerant algorithms and data structures. Furthermore, *Henrik Blunck* (PhD Munster 2006) continues to be affiliated with the center even though he is now working on the Galileo project funded by the high-tech foundation (Højteknologifonden). Note that one of the 2009 hires (Sitchinava) visited the center for two extended periods while a PhD student at UC Irvine.

PhD students:

- *Martin Olsen* (Advisor: Brodal).
  Worked on algorithms for analyzing the web-graph; defended on November 9, 2009.
  Spent 6 months at ICT, Sidney. Now faculty at AU Institute of Business and Technology, Herning.
- *Thomas Mølhave* (Advisor: Arge).
  Worked on I/O-efficient terrain data and fault tolerant algorithms; defended on October 21, 2009.
  Spent 2 months at AT&T research and 4 months at Duke University. Now Post Doc at Duke.
- *Allan G. Jørgensen* (Advisor: Brodal).
  Worked on data structures and fault tolerant algorithms; handed in thesis in January 2010.
  Spent the spring of 2008 at MIT.
- *Lasse Deleuran* (Advisor: Arge).
  Works on line simplification (computational geometry) algorithms; will finish PhD in July 2011.
  Is spending the spring of 2010 at Duke University.
- *Kostas Tsakalidis* (Advisor: Brodal).
  Works on I/O-efficient interpolation search; will finish PhD in July 2011.
  Will spend a semester at MIT and University of Waterloo.
- *Mark Greve* (Advisor: Brodal).
  Works on I/O-efficient sorted range reporting; will finish PhD in January 2012.
- *Morten Revsbæk* (Advisor: Arge).
  Works on I/O-efficient terrain data algorithms; will finish PhD in January 2013.
  Is on leave-of-absence between November 2009 and October 2010 to work at SCALGO.
- *Jesper E. Moeslund* (tidligere Eshøj) (Advisor: Arge with Biology and Agricultural Sciences faculty)
  Works on plant diversity modeling; will finish PhD in January 2012.
  Will spend the fall semester at ETH-WSL.
- *Pooya Davoodi* (Advisor: Brodal).
  Works on range minimum data structures; will finish PhD in April 2011.
  Will spend a semester at Leicester and Carleton Universities.
- *Jacob Truelsen* (Advisor: Brodal).
  Works on range mode structure lower bounds; will finish PhD in July 2012.
- *Kasper D. Larsen* (Advisor: Arge).

Works on geometric data structures; will finish PhD in July 2013.
- *Casper Kejlberg-Rasmussen* (Advisor: Brodal).
  Started in February 2009; will finish PhD in July 2013.
- *Freek van Walderveen* (Advisor: Arge).
  Started in August 2009; will finish PhD in July 2012.
- *Sarah Zakarias* (Honors master student - has yet to choose PhD topic/advisor)
  Started in August 2009; will finish PhD in July 2014.

In addition to the above students, *Andreas Beckmann, Andrei Negoescu* and *Volker Weichert* are center PhD students at FRA (with Meyer as advisor), and *Kahn Do Ba* and *Eric Price* (with Indyk as advisor), and *Jelain Nelsen* and *Morteza Zadimoghaddam* (with Demaine as advisor) are center PhD students at MIT. During 2009 center PhD student *Oren Weimann* (Advisor: Demaine) obtained his PhD from MIT. He is now a Post Doc at the Weizmann Institute.

As also discussed in the 2008 annual report, we believe the centers Post Doc and PhD student recruitment efforts have been relatively successful, although we would like to recruit more international PhD students (so far we have recruited one international student at AU in each of the three years of the center). In general, the hiring of PhD students has progressed more or less as planned. In fact, the number AU center PhD students is now higher than the total number anticipated in the center contract. Still, depending on the strength of the applicants, new AU PhD students might be affiliated with the center this summer (funded by other sources). Currently faculty resources is a major constraint on PhD student admission at AU, but hopefully this constrain should be removed when the open Associate Professor position is filled this fall.
As mentioned, the centers focus on research education includes exchange of Post Doc and PhD students, as well as a 6 months stay abroad for AU PhD students. Previously, the center has also organized two international summer schools, but in 2009 focus was on organizing the Symposium on Computational Geometry and the Workshop on Massive Data Algorithmics. However, we are planning another summer school in August 2010. Center Post Docs also continue to organize specialized PhD classes; in 2009 a 5 EATCS class on Advanced Randomized Algorithms and in 2010 classes on I/O-Efficient Graph Algorithms and on Advanced Geometric Data Structures. Finally, the center continues to emphasize initiatives designed to create as sense of community at the main center site and among the center sites. This includes a yearly two day fall retreat, monthly center lunches at AU, as well as number of social events (Christmas dinner, weekly breakfast, weekly soccer, go-carting trips, etc).

6 Milestones

As discussed in the research plan, it is often very hard to establish very concrete milestones and goals for the kind of theoretical research conducted in the center. In Section 2, we have discussed results obtained in 2009 and their relation to the research plan, as well as outlined a research agenda for 2010 (and beyond).

A significant research production (published in the major conference and journals) and high visibility in the research community is obviously one of the major goals of the center. We believe both of these goals are being meet. The 2008 annual report also identified a number of high-level 2009 milestones. Most of these milestones have been met: Organization of the first Workshop on Massive Data Algorithmics and the 25th Annual ACM Symposium on Computational Geometry, further work on commercialization of the TerraSTREAM, development of further PhD classes, as well as a renewed focus on collaborative projects.
We still have to organize a larger multidisciplinary workshop on terrain data handling and analysis, and we are still considering the establishment of video conference and remote collaboration facilities between the centers cites, as well as the establishment of more formalized international PhD student and summer student programs. We have decided not to establish a Danish advisory board (as a supplement to our international advisory board). We also still plan on developing educational material on I/O-efficient algorithms, but are reluctant to set it as a milestone for 2010. Other milestones for 2010 include organization of the second Workshop on Massive Data Algorithmics and of an international summer school (probably on advanced geometric data structures). We are also planning a couple of new events either in 2010 or early 2011, namely an invitation only workshop at massive data algorithms at Sandbjerg (using a format that is well-known among computer scientists from the German “Dagstuhl” meetings), as well as an interdisciplinary Symposium on “Biodiversity in the Silicon Age”.

Appendix to MADLGO Annual report 2009

Appendix A: External Relations
Please list collaborations with danish and foreign research groups, universities, companies and/or institutions and state subject and/or title and output/results if any.

List **includes** relevant 2009 collaboration for AU as well as MIT, MPI and FRA researchers.

			Collaboration with: (Please tick the appropriate box)			
Collaborator Name (person and/or institution) Country	**Collaboration subject and/or title**	**Output/ results of collaboration, if any**	**Danish universities, research groups and institutions**	**Foreign universities, research groups and institutions**	**Danish companies**	**Foreign companies**
Peder Klith Bøcher (Faculty of Agricultural Science, AU), Jens-Christian Svenning (Institute of Biological Sciences, AU), and National Environmental Research Institute (NERI), Denmark	Collaborators and co-advisors of PhD student Jesper Moeslund Eshøj ("Fintopløselig, geospatial modellering af Danmarks nuværende og potentielle fremtidige plantediversitet")	Publication	X			
Pankaj K. Agarwal and Bardia Sadri (Duke), Andrew Danner (Swarthmore College), USA	I/O-efficient terrain algorithms and TerraSTREAM	Publications, as well as the TerraSTREAM software package		X		
Timothy Chan (Waterloo), Canada, and Jeremy Barbay (Universidad de Chile), Chile	Geometric algorithms and data structures	Publications		X		
BNR A/S, Denmark	GIS in traffic management	Software			X	
Hamish Carr (Leeds), England	Scalar field simplification			X		
Herman Haverkort (TU Eindhoven), the Netherlands	I/O-efficient removal of noise from terrain data	MSc Thesis		X		
Norbert Zeh (Dalhausie), Canada	I/O-efficeint, cache-oblivious and parallel algorithms	Publications		X		
COWI A/S (incl. Kristian Keller, Johnny Koust Rasmussen, Michael Schultz Rasmussen) and The Faculty of Agricultural Science (incl. Peder Klith Bøcher), Denmark	Efficient Handling of Massive Heterogenious Terrain Data	Efficient terrain processing algorithms and TerraSTREAM software modules	X		X	
Jan Vahrenhold (TU Dortmund), Germany and Andrew Danner (Swarthmore College), USA	TPIE	TPIE software package		X		
Mike Goodrich (UC Irvine), USA	Algorithms for private-cache chip multiprocessors	Publications		X		

Ian Munro and Alex Lopez-Ortiz (Waterloo), Norbert Zeh (Dalhousie), Bradford G. Nickerson (New Brunswick), Mark Masry (CARIS Geomatics Software Solutions), Canada	Data Structures for Efficient Organization and Retrieval of Massive Spatial Data			X		X
Eiva A/S, Denmark	Sonar data cleaning	Publication, software				X
Mark de Berg and Herman Haverkort (TU Eindhoven), Netherlands	Geometric algorithms and data structures	Publications		X		
Researcher at Aarhus and Aalborg University, Terma A/S, Systematic Software Engineering A/S, Dansk Landbrugsrådgivning, Alexandra Instititute, Denmark	Hightech foundation project "A platform for Galileo based pervasive computing"	Publications	X		X	
Scalable Algorithmics (SCALGO), Denmark	I/O-efficient terrain algorithms and software				X	
Kostas Tsichlas, Apostolos N. Papadopoulos (Thessaloniki), Spyros Sioutas (Ionian), Alexis Kaporis (Patras), Greece	Range searching	Publications		X		
Rolf Fagerberg (University of Southern Denmark), Denmark, Alex López-Ortiz (Waterloo), Canada	Sorted Range Reporting	Publication	X	X		
John Iacono (NYU), Stefan Lagnerman (Bruxelles), Ian Munro (Waterloo), USA, Belgium and Canada	Cache-Oblivious Dictionaries	Publication		X		
Xavier Freixas, Xavier Molinero Albareda, Maria Serna (Barcelona), Spain	Simple Games	Manuscript		X		
Taso Viglas (Sydney), Australia	Link Building	Manuscript		X		
Srinivasa Rao (Seoul), South Korea, Rajeev Raman (Leicester), UK	Range Minimum Query	Manuscript		X		
Riko Jacob (TU Muenchen), Germany	Computational Models for Flash Memory Devices	Publication		X		
Mark de Berg (TU Eindhoven) The Netherlands, Mohammad Farshi, Michiel Smid (Carleton), Canada, Joachim Gudmundsson (NICTA), Australia	Geometric spanners	Publications		X		

Mark de Berg (TU Eindhoven), The Netherlands, Alireza Zarei (Sharif), Iran	Streaming Algorithms for Line Simplification	Publication		X		
Mohammad Farshi, Michiel Smid (Carleton), Canada, Paz Carmi (Ben-Gurion), Israel	Semi-seperated Pair Decompositions	Publication		X		
Rasmus Pagh (IT University Copenhagen), Denmark	Secondary Indexing in One Dimension	Publication	X			
Roberto Grossi, Alessio Orlandi (University of Pisa), Italy, Rajeev Raman (University of Leicester), Great Britain	Redundancy in Fully Indexable Dictionarie	Publication		X		
Jyrki Katajainen (University of Copenhagen), Denmark	Dynamic Dictionaries for Multisets	Manuscript	X			
Rajeev Raman (Leicester), Great Britain, Arash Farzan (Waterloo), Canada	Succinct representations of trees	Publication		X		
Martin Smerek (Masaryk University), Czech Republic	I/O-efficient software for BDD manipulation	Software		X		
Maarten Löffler (UC Irvine), USA, Jeff M. Phillips (Utah), USA	Geometric Computations on Indecisive Points	Manuscript		X		
Henning Meyerhenke (Paderborn), Germany	Realistic Computer Models	Publication		X		
Jan Vahrenhold, Fabian Gierseke (TU Dortmund), Germany	Resilient kd-kmeans			X		
Lars Engebretsen, Google Zurich	Graph Clustering					X
German Algorithm Engineering Cluster	Selected Topics in Alg. Engineering	Workshops		X		
Lufthansa Systems, Frankfurt, Germany	Efficient shortest-paths computations with dynamic weights					X
Group of Peter Sanders (Karlsruhe), Germany	(1) Libraries for parallel/external computation (2) Implementation of EM SSSP (3) Energy-efficient sorting	Publications		X		
David Woodruff (IBM Almaden), USA	Streaming algorithms	Papers				X
Rich Baraniuk (Rice), USA	Compressive sensing	Paper		X		
Atri Rudra (SUNY), USA	Group testing	Paper		X		
Graham Cormode (AT&T), Martin Strauss (Michigan), USA	Heavy hitters	Paper				X
Robert Krauthgamer (Weizman), Israel	Embeddings	Papers		X		

John Iacono (Polytechnic Inst. NYU), Hayley Iben and James O'Brien (Berkeley), Goran Konjevod (Arizona State), Vi Hart (SUNY Stony Brook), Joseph O'Rourke (Smith College), and Robert J. Lang (Lang Origami), USA, Jean Cardinal and Stefan Langerman (Bruxelles), Belgium, Shinji Imahori (U. Tokyo) and Ryuhei Uehara (JAIST), Japan, Sandor Fekete (TU Braunschweig), Germany	Folding	Publications		X		X
Mihai Badoiu (Google), Martin Farach-Colton (Rutgers) and MohammadTaghi Hajiaghayi (AT&T Research), USA, Anastasios Sidiropoulos (Toronto), Canada, Noga Alon (Tel Aviv), Israel	Ordinal Embeddings	Publication		X		X
Francisco Gomez-Martin (U. Politecnica de Madrid), Spain, Henk Meijer and David Rappaport (Queens), Godfried Toussaint and David Wood (McGill), Canada, Perouz Taslakian (Bruxelles), Belgium, Terry Winograd (Stanford), USA	Music Distance Geometry	Publication		X		
Greg Aloupis (McGill) and Stefanie Wuhrer (Carleton), Canada, Sebastien Collette and Stefan Langerman (Bruxelles), Belgium, Robin Flatland (Siena College), Joseph O'Rourke (Smith College), James McLurkin (Rice), Suneeta Ramaswami (Rutgers), Mirela Damian (Villanova), and John Iacono (Polytechnic Inst. NYU), USA, Vera Sacristan (U. Politecnica de Catalunya), Spain	Modular Robots	Publications		X		
Stefan Langerman (Bruxelles), Belgium, Gad M. Landau (Haifa), Israel, Dion Harmon (New England Complex Systems Institute), John Iacono (Polytechnic Inst. NYU), USA	Data structures	Publications		X		X

Jonathan Bredin (Colorado College), MohammadTaghi Hajiaghayi (AT&T Research), USA	Sensor Networks	Publication		X		X
MohammadTaghi Hajiaghayi (AT&T Research), USA, Hamid Mahini, Amin S. Sayedi-Roshkhar and Shayan Oveisgharan (Sharif U. Technology), Iran, Daniel Marx (Tel Aviv U.), Israel	Game theory and approximation algorithms	Publications		X		X
Takehiro Ito (Tohoku), Japan, Marcin Kaminski (Bruxelles), Belgium, Brad Ballinger (Davis), David Charlton (Boston), John Iacono (Polytechnic Inst. NYU), MohammadTaghi Hajiaghayi (AT&T Research), Philip Klein and Shay Mozes (Brown), USA, Glencora Borradaile (Waterloo), Canada, Ching-Hao Liu and Sheung-Hung Poon (National Tsing Hua U.), China, Siamak Tazari (Berlin), Germany, Raphael Yuster and Ilan Newman (Haifa), Israel, Jean Cardinal, Samuel Fiorini and Gwenael Joret (Bruxelles), Belgium	Graph algorithms	Publications		X		
Sandor Fekete and Nils Schweer (TU Braunschweig), Gunter Rote and Daria Schymura (FU Berlin), Germany	Integer Point Sets Minimizing Average Pairwise L_1 Distance	Publication		X		
Robert A. Hearn (Dartmouth), USA	Games, Puzzles, and Computation	Book		X		
Michael A. Burr (NYU), John Hugg and Kathryn Seyboth (Tufts), Eynat Rafalin (Google), USA, Stefan Langerman (Bruxelles), Belgium, Timothy M. Chan (Waterloo), Canada	Dynamic Ham-Sandwich Cuts in the Plane	Publication		X		X
Miklos Ajtai and Vitaly Feldman (IBM Research), USA	Sorting and Selection with Imprecise Comparisons	Publication		X		X
Rolf Backofen (Freiberg) and Gad M. Mathias Möhl (Saarland), Germany, Dekel Tsur (Ben-Gurion) and Landau (Haifa), Israel	Fast RNA Structure Alignment for Crossing Input Structures	Publication		X		

Appendix B: Conferences

Please outline the centres disseminations of results and networking through participation in large internationale conferences, symposia etc. by a) list the number of conferences, symposia, seminars etc. the centre has arranged or participated in the planning of and by b) list the number of conferences, symposia, seminars etc. the centre has participated in with contributions.

List **includes** relevant 2009 information (conference participation with paper/poster or as PC member/organizer) for AU as well as MIT, MPI and FRA researchers.

a) Organisation of international conferences, symposia, seminars etc.

Title of event	**Number of participants, estimated**	
	Danish	**International**
Workshop on Massive Data Algorithmics (MASSIVE)	25	38
Symposium on Computational Geometry (SoCG)	26	110
25th Symposium on Computational Geometry Celebration	20	98

b) Participation in international conferences, symposia, seminars etc.

Title of event	**Venue**	**Name(s) of participant(s)**	**Contribution (talk, abstract, paper, poster, other)**	**Invited talk (please check)**
Beyond Kyoto: Addressing the Challenges of Climate Change. Science meets Industry, Policy and Public	Aarhus University, Denmark	Moeslund, Arge	Poster	
Mapping and Monitoring of Nordic Vegetation and Landscapes	Hveragerði, Iceland	Moeslund	Poster	
International Symposium on Experimental Algorithms	Dortmund, Germany	Hachenberger, Beckmann	Talks	
Symposium on Foundations of Computer Science	Atlanta, USA	Larsen, Afshani, Arge, Do Ba	Talks	
Algorithms and Data Structures Symposium	Banff, Canada	Mølhave, Demaine	Talks	
Annual Symposium on Computational Geometry	Aarhus, Denmark	Arge, Brodal, Indyk, Demaine, Mehlhorn, Afshani, Abam (and others)	Talks, Organizing	

Current Trends in Algorithms, Complexity Theory, and Cryptography	Beijing, China	Afshani, Arge, Brodal	Talks	X
Workshop on Theory and Many-Cores (T&MC)	College park, MD, USA	Sitchinava	Talk	
Seminar	Cambridge University, GB	Sitchinava	Talk	X
International Symposium on Algorithms and Computation	Honolulu, USA	Revsbæk, Arge, Greve, Jørgensen, Tsakalidis, Negoescu, Moruz, Demaine	Talks	
European Workshop on Computational Geometry	Brussels, Belgium	van Walderveen	Paper	
European Symposium of Algorithms	Copenhagen, Denmark	Arge, Brodal, Abam, Mehlhorn, Meyer, Demaine	Steering Committee, Talks	X
International Symposium on Spatial and Temporal Databases	Aalborg, Denmark	Arge	Talk, PC	X
Symposium on Principles of Database Systems	Providence, RI, USA	Arge	Talk	X
Seminar	Duke University, USA	Afshani	Talk	X
Seminar	IMADA, SDU, Odense	Greve	Talk	X
Seminar	ITU, Copenhagen	Greve	Talk	X
Workshop on Massive Data Algorithmics	Aarhus, Denmark	Arge, Brodal, Jørgensen, Ajwani, Beckmann (and others)	Talks, Organizing	
Dagstuhl seminar on Computational Geoemtry	Dagstuhl, Germany	Abam, Mehlhorn, Arge	Talks	X
Dagstuhl seminar on Geometric Networks, Metric Space Embeddings and Spatial Data Mining	Dagstuhl, Germany	Abam	Talk	X
Seminar	Institute of Mathematical Sciences, Chennai, India	Rao	Talk	X
Dagstuhl seminar on Graph Search Engineering	Dagstuhl, Germany	Ajwani, Meyer	Talks	X
Workshop on Algorithm Engineering and Experiments	New York, USA	Meyer	Paper	

International Symposium on Parallel and Distributed Processing	Miami, FL, USA	Beckmann	Paper	
Int. Conference on Symbolic and Algebraic Computation	Seoul, Korea	Mehlhorn	Paper	
Seminar of the German Research Cluster Algorithm Engineering	Dortmund, Germany	Moruz	Talk	X
Kurt Mehlhorn 60th Birthday Colloquium	Saarbrücken, Germany	Meyer	Talk	X
Colloquium on Automata, Languages and Programming	Rhodes, Greece	Mehlhorn, Demaine, Nelson	Talks	X
Seminar	Dortmund, Germany	Moruz	Talk	X
Information Theory and Applications Workshop	San Diego, USA	Indyk	Talk	X
Seminar	Rice University, USA	Indyk	Talk	X
EPFL Summer Research Institute	Lausanne, Switzerland	Indyk	Talk	X
Conference on Learning Theory	Montreal, Canada	Indyk	Talk	X
Workshop on Algorithms for Data Streams	IIT Kanpur, India	Indyk, Nelson	Talks	X
New Directions in Applied Mathematics Workshop	IISc, Bangalore	Indyk	Talk	X
Microsoft Research-India Theory Day	Chennai, India	Indyk	Talk	X
Complexity reading group	Microsoft Research New England, USA	Nelson	Talk	X
Seminar	University of Michigan, USA	Nelson	Talk	X
China Theory Week	Beijing, China	Nelson	Talk	X
Seminar	California Institute of Technology, USA	Nelson	Talk	X
Seminar	University of California, Los Angeles, USA	Nelson	Talk	X
Seminar	IBM Almaden Research Center, USA	Nelson	Talk	X
Seminar	IT University of Copenhagen, Denmark	Nelson	Talk	X

DIMACS/DyDan Workshop on Streaming, Coding, and Compressive Sensing	DIMACS, New Brunswick, USA	Nelson	Talk	X
Conference of BRIDGES: Mathematical Connections in Art, Music, and Science	Banff, Canada	Demaine	Talk	X
Workshop on Algorithmic Aspects of Wireless Sensor Networks	Rhodes, Greece	Demaine	Talk	X
Japan Conference on Computational Geometry and Graphs	Kanazawa, Japan	Demaine	Talks, Papers	X
British Science Festival	Guildford, England	Demaine	Talk	X
British Origami Society Autumn Convention	Winchester, England	Demaine	Talk	X
International Fab Lab Forum and Symposium on Digital Fabrication	India	Demaine	Talk	X
Annual Meeting of the American Association for Advancement of Science	Chicago, IL, USA	Demaine	Talk	X
Workshop on Internet & Network Economics	Rome, Italy	Demaine	Paper	
IEEE/RSJ International Conference on Intelligent Robots and Systems	St. Louis, Missouri, USA	Demaine	Paper	
European Conference on Mobile Robots	Mlini/Dubrovnik, Croatia	Demaine	Paper	
Canadian Conference on Computational Geometry	Vancouver, Canada	Demaine	Papers	
European Workshop on Computational Geometry	Brussels, Belgium	Demaine	Paper	
International Symposium on Theoretical Aspects of Computer Science	Freiburg, Germany	Demaine	Talk, Papers	
ACM-SIAM Symposium on Discrete Algorithms	New York, NY, USA	Demaine	Talk, Papers	
Seminars	Vrije University, Universite Catholique de Louvain, Universite Libre de Bruxelles, Brussels, Belgium	Demaine	Talk	X
Seminar	Northwestern University, Chicago, IL, USA	Demaine	Talk	X

Appendix C: Educational activities

Please list all educational activities the centre is envolved in. Including PhD-courses, courses on bachelor- and master-level and other activities on academic level as well as summer schools and courses teached abroad. Please state ECTS points (if possible) and length of the course (in hours)

List only **includes** 2009 information for AU employees (as well as relevant information for MIT, MPI and FRA employees **taught outside** their home institution). PhD student TA'ing is **not** included.

Title and date of activity	**ECTS**	**Length of course (number of hours)**	**Number of participants**
BSc course: Computer Science in Perspective, Fall 2009	5	10	120
BSc course: Algorithms and Data Strucutres 1, Spring 2009	5	28	170
BSc course: Algorithms and Data Strucutres 2. Spring 2009	5	28	135
BSc course: Basic Software Development 1, Institute of Business and Technology, Fall 2009	5		15
BSc course: Basic Software Development 2, Institute of Business and Technology, Spring 2009	5		15
MSc course: I/O-efficient algorithms, spring 2009	10	36	12
MSc course: Advanced Data Structures, Fall 2009	10	42	40
MSc Honors class: Advanced Data Structures, Fall 2009	10	42	10
PhD course: Advanced Randomized Algoritihms, Spring 2009	5	21	10
Invited tutorial on Worst-Case Efficient Range Search Indexing, ACM Symposium on Principles of Database Systems, Fall 2009		2	~30
DM Programming Contest Coaching			20
NWERC Programming Contest Coaching			6
Graduate level course: Streaming etc, Rice University, Spring 2009		9	20
Summer research institute tutorial: Streaming, Sketching and Sub-linear Space Algorithms, EPFL, Summer 2009		1	~30

Students surpervised at the centre by staff members (AU, MPI and FRA)

Number of graduated PhD-students	**Number of graduated Master-students**
3	4

Appendix D: Academic services

Please list the activities for the entire group of centre staff members.

List **only includes** 2009 information for AU employees.
Conference reviewing is also included since conferences play a large role in computer science; the lists are **probably incomplete** and numbers **only rough estimates**, since computer scientists typically do a very large number of (especially conference) reviews each year.

a) Scholarly communication

Title of the journal	**Editorial work (please check)**	**Number of peer reviews and pre-assess-ments done for the journal**
Journal of Computational Geometry	X	0
Journal of Experimental Algorithmics	X	3
Journal of Graph Algorithms and Applications	X	2
Journal of Discrete Algorithms	X	0
Algorithmica	X	7
Electronic Proceedings in Theoretical Computer Science	X	1
International Journal of Computational Geometry and Applications		3
Journal on Computers & Graphics		2
ACM Transactions on Graphics		1
Journal on Information Sciences		1
Information Processing Letters		2
Transactions on Algorithms		2
Journal of ACM		1
Algorithmica		3
Logical Methods in Computer Science		1
Information and Computation		1
Workshop on Massive Data Algorithmics	X	48
ACM International Symposium on Advances in Geographic Information Systems	X	~7
International Symposium on Spatial and Temporal Databases	X	~6
International Computing and Combinatorics Conference	X	~17
International Symposium on Experimental Algorithms	X	16
ACM Symposium on Parallelism in Algorithms and Architectures	X	22
Workshop on Algorithms and Data Structures	X	10
Latin American Theoretical Informatics Symposium	X	12
Other reviewing for major conferences (where not on Program Committee)		~25

b) Assessments, international panels, membership and supervision

Number of staff conducting academic services	**Number of grant and/or career/tenure assessments plus international panels**	**Number of positions and council seats in learned societies**	**Number of Ph.d. and doctor evaluations, national and international**	**Number of primary supervisions of Ph.D and master students**
2	2	3 conference steering committees, one membership of academy of science and letters	2	13PhD, 4 MS

Appendix E: External funding

Please list all external funding gained by the centreleader or centre members. List only funding gained for centre activities including external PhD-grants. List the full amount and the part of the total amount allocated to reported year.

List **only includes** information for AU employees; it includes all active/new funding in 2009

	Funding body	**Purpose**	**Applicant**	**Activity period**	**Granted amount in DKK**	**Part of total amount allocated to reported**
Public Danish funds	Danish National Advanced Technology Foundation (Højteknologi-fonden)	A platform for Galileo based pervasive positioning	AU (incl Arge), AAU, Danish agricultural advisory service and several companies including Alexandra A/S, Terma A/S, and Systematic	2007-2010	~11.700.000 (AU part)	?
	Strategic Research Council (NABIIT program)	Efficient Handling of Massive Heterogeneous Terrain Data	AU (Arge), DJF and COWI A/S	2006-2010	~7.200.000	~2.000.000
	Forsknings- og innovationsstyrelsen, PhD kvalitetsfremmemidler	Post Doc support	Arge (through BRICS)	2009	~250.000	~250.000
Private Danish funds	Aarhus Universitets Forskningsfond	Center for Interdisciplinary Geospatial Informatics Research	Faculty of Natural Sciences, AU (incl Arge)	2009-2012	2.500.000	0
	VELUX Visiting Professor Program	Visit of Norbert Zeh	Arge (through NF, AU)	2009	80.000	80.000
International funds	US Army Research Office	STREAM: Scalable Techniques for High Resolution Elevation Data Analysis and Modeling	Duke University, NCSU, and AU (Arge)	2009-2011	~2.800.000	?
	NORFA	Nordic Network on Algorithms	Brodal (through Bergen)	2008-2009	510000	255000

Appendix F: Awards

Please list awards or other acknowledgements gained by centre members in reported year.

List **includes** relevant 2009 information for AU as well as MIT, MPI and FRA researchers.

Awards	**Recipient**	**Granted amount in DKK, if relevant**
Elected Distinguished Scientist by Association for Computing Machinery (ACM)	Arge	
George Polya Lecturer, Mathematical Association of America	Demaine	
Appointed Associate Editor of ACM Journal on Experimental Algorithmics	Arge	
Appointed Associate Professor with Exceptional Qualifications (Lektor MSK)	Brodal	

Appendix G: Public outreach

Please list public outreach activities in media, press, hign schools etc.

List **only includes** 2009 information for AU employees.

a)Electronic media

Specific media (TV, radio, other)	**Date**	**Type of communication (interview, commentary, debate, feature program etc)**	**Subject/ Title**	**Contributor from the centre**
videnskab.dk	11. marts 2009	Feature	Dansk sofware skal varsle oversvømmelser	Arge
Computerworld	12. marts 2009	Feature	Ny software holder danskerne tørskoede	Arge
TV2.dk	12. marts 2009	Feature	Ny software holder danskerne tørskoede	Arge
Version2	18. september 2009	Feature	AU-forskere satte turbo på GIS-beregninger med ram-optimering	Arge
Ingeniøren.dk	18. september 2009	Feature	Dansk algoritme sætter turbo på digitale terrænmodeller	Arge

b)Press

Specific media (Daily newspapers, journals, magazines, other)	**Date**	**Type of communication (inter-view, commen-tary, debate, feature etc)**	**Subject/ Title**	**Contributor from the centre**
Jyllandsposten	5. april 2009	Feature	Ny software varsler oversvømmelser	
Ingeniøren	18. september 2009	Feature	Dansk algoritme sætter turbo på digitale terrænmodeller	Arge
Weekendavisen	2. oktober 2009	Feature	Lysindfald ganget med vand	Arge
Aarhus Stiftstidende	16. december 2009	Front page	Klimaændringer vil drukne Århus	Moeslund
Aarhus Stiftstidende	16. december 2009	Interview	Ny målemetode tegner et nøjagtigt højdekort	Moeslund
Aarhus Stiftstidende	16. december 2009	Feature	Klimaændringer slår hårdt i Risskov og Egå	Moeslund
Geoforum	1. maj 2010	Article	GIS ved MADALGO	Arge, Mølhave

c)Other (talks at secondary educational institutions etc)

Specific type of communication (presentation/lecturing at open university, high school etc)	Date	Subject/Title	Contributor from the centre
Presentation at "Kortdage 2009"	19. november 2009	Hvor løber vandet hen? Oversvømmelsesberegninger på store højdemodeller	Arge
Presentation at "COWI GIS konference 2009"	23. september 2009	Oversvømmelse og vandstrømning	Arge
Presentation at "Datalogforeningen - faglig aften"	27. maj 2009	MASSIVE Data Algorithmics	Arge
Presentation at "Rotary - Ebeltoft"	27. april 2009	Hvornår bliver rådhuset oversvømmet - og hvor kommer vandet fra?	Arge
Presentation at "Geoforum højdedataseminar"	11. marts 2009	Algorithms for Handling Massive (Elevation) Datasets	Mølhave
Lectures at ScienceTalenter Center, Sorø for Masterclass in Mathematics (Selected High School students).	October 29-31, 2009	Algorithms: Matrices and Graphs	Brodal
Presentation at "Ungdommens Naturvidenskabelige Forening, Ålborg"	September 29, 2009	Internet Search Engines	Brodal
Presentation at MasterClass Coordination Meeting for a group of High School teachers	September 4, 2009	MADALGO	Brodal
Gymnasie Studiepraktik	October 28, 2009	Lecture/Excerises on algorithms	Brodal

Appendix H: Patents and applications

List the number of submitted patent applications, granted patents etc. gained by the centre in 2009. Also list possible spin-off companies and collaborations/partners. etc.

List **only includes** 2009 information for AU employees.

Number of submitted patent applications	**Number of granted patents**	**Number of mutually agreed licence, sale and option agreements**	**Names of spin-off companies established**	**Other application (please specify)**	**Partner (s), if any**
			Scalable Algorithmics (SCALGO)		COWI A/S, EIVA A/S

Appendix I: Publications

Please enclose, in a seperate appendix, a full publication list including all authors dating back from the centre start. The publication list shall be arranged in cronological order and numbered so the most recent publication gets the highest number and is placed at the bottom of the list. Divide the list in types of publications including conference proceedings/presentations as well as master- and PhD-thesis made at the centre. List only accepted publications and don't include publications that are "submitted" or "in press". List only publications authored or co-authored by centre members. For each publication state number, author(s), place and year of publication. If the publication have been peer reviewed please note in brackets "PR". Also note in brackets "CO" if the publiation is co-authored by another national or international research institution.

List **includes** relevant publications for AU as well as MIT, MPI and FRA researchers

Conference proceedings

1	2007	B. Escoffier, G. Moruz and A. Ribichini	Adapting Parallel Algorithms to the W-Stream Model, with Applications to Graph Problems	Proc. International Symposium on Mathematical Foundations of Computer Science (MFCS)	(PR)(CO)
2	2007	S. Guha, P. Indyk and A. McGregor	Sketching Information Divergences	Proc. Annual Conference on Learning Theory (COLT)	(PR)(CO)
3	2007	G. S. Brodal and A. G. Jørgensen	A Linear Time Algorithm for the k Maximal Sums Problem	Proc. International Symposium on Mathematical Foundations of Computer Science (MFCS)	(PR)(CO)
4	2007	G. S. Brodal, L. Georgiadis, K. A. Hansen and I. Katriel	Dynamic Matchings in Convex Bipartite Graphs	Proc. International Symposium on Mathematical Foundations of Computer Science (MFCS)	(PR)(CO)
5	2007	G. Jørgensen, G. Moruz and T. Mølhave	Resilient Priority Queues	Proc. International Workshop on Algorithms and Data Structures (WADS)	(PR)
6	2007	G. S. Brodal, R. Fagerberg, I. Finocchi, F. Grandoni, G. Italiano, A. G. Jørgensen, G. Moruz and T. Mølhave	Optimal Resilient Dynamic Dictionaries	Proc. European Symposium on Algorithms (ESA)	(PR)(CO)

7	2007	P. K. Agarwal, L. Arge, A. Danner, H. Mitasova, T. Mølhave and K. Yi	TerraStream: From Elevation Data to Watershed Hierarchies	Proc. ACM International Symposium on Advances in Geographical Information Systems (ACM-GIS)	(PR)(CO)
8	2007	M. Patrascu and Mikkel Thorup	Planning for Fast Connectivity Updates	Proc. IEEE Symposium on Foundations of Computer Science (FOCS)	(PR)(CO)
9	2007	G. Franceschini, S. Muthukrishnan, and M. Patrascu	Radix Sorting With No Extra Space	Proc. European Symposium on Algorithms (ESA)	(PR)(CO)
10	2007	E. D. Demaine, S. Mozes, B. Rossman and O. Weimann	An Optimal Decomposition Algorithm for Tree Edit Distance	Proc. International Colloquium on Automata, Languages and Programming (ICALP)	(PR)(CO)
11	2007	M. A. Bender, M. Farach-Colton, J. T. Fineman, Y. Fogel, B. C. Kuszmaul and J. Nelson	Cache-Oblivious Streaming B-trees	Proc. ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)	(PR)(CO)
12	2007	E. D. Demaine, M. Ghodsi, M. Hajiaghayi, A. S. Sayedi-Roshkhar and M. Zadimoghaddam	Scheduling to Minimize Gaps and Power Consumption	Proc. ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)	(PR)(CO)
13	2007	M. Patrascu	Lower Bounds for 2-Dimensional Range Counting	Proc. ACM Symposium on Theory of Computing (STOC)	(PR)
14	2007	G. M. Landau, D. Tsur and O. Weimann	Indexing a Dictionary for Subset Matching Queries	Proc. Symposium on String Processing and Information Retrieval (SPIRE)	(PR)(CO)
15	2007	T. Friedrich and D. Ajwani	Average-Case Analysis of Online Topological Ordering	Proc. International Symposium on Algorithms and Computation (ISAAC)	(PR)
16	2007	K. Chang	Multiple pass streaming algorithms for learning mixtures of distributions in R^d	Proc. Algorithmic Learning Theory (ALT)	(PR)

17	2007	M. Westergaard, L. M. Kristensen, G. S. Brodal and L. Arge	The ComBack Method - Extending Hash Compaction with Backtracking	Proc. International Conference on Applications and Theory of Petri Nets and Other Models of Concurrency (ICATPN)	(PR)
18	2007	M. A. Bender, G. S. Brodal, R. Fagerberg, R. Jacob and E. Vicari	Optimal Sparse Matrix Dense Vector Multiplication in the I/O-Model	Proc. ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)	(PR)(CO)
19	2007	A. Golynski, R. Grossi, A. Gupta, R. Raman and S. S. Rao	On the Size of Succinct Indices	Proc. European Symposium on Algorithms (ESA)	(PR)(CO)
20	2007	M. Olsen	Nash Stability in Additively Separable Hedonic Games is NP-hard	Proc. Conference on Computability in Europe (CiE)	(PR)
21	2008	M. Ruzic and P. Indyk	Near-Optimal Sparse Recovery in the L1 norm	Proc. Symposium on Foundations of Computer Science (FOCS)	(PR)(CO)
22	2008	M. Patrascu	(Data) STRUCTURES	Proc. Symposium on Foundations of Computer Science (FOCS)	(PR)
23	2008	M. Patrascu	Succincter	Proc. Symposium on Foundations of Computer Science (FOCS)	(PR)
24	2008	E. Demaine, S. Langerman and E. Price	Confluently Persistent Tries for Efficient Version Control	Proc. Scandinavian Workshop on Algorithm Theory (SWAT)	(PR)(CO)
25	2008	D. Ajwani, I. Malinger, U. Meyer and S. Toledo	Characterizing the Performance of Flash Memory Storage Devices and Its Impact on Algorithm Design	Proc. Workshop on Experimental Algorithms (WEA)	(PR)(CO)
26	2008	U. Meyer	On Dynamic Breadth-First Search in External-Memory	Proc. Symposium on Theoretical Aspects (STACS)	(PR)
27	2008	U. Meyer	On Trade-Offs in External-Memory Diameter Approximation	Proc. Scandinavian Workshop on Algorithm Theory (SWAT)	(PR)

28	2008	G. S. Brodal and A. G. Jørgensen	Selecting Sums in Arrays	Proc.International Symposium on Algorithms and Computation (ISAAC)	(PR)
29	2008	L. Arge, G. S. Brodal and S. S. Rao	External Memory Planar Point Location with Logarithmic Updates	Proc. Symposium on Computational Geometry (SoCG)	(PR)
30	2008	A. Golynski, R. Raman and S. S. Rao	On the Redundancy of Succinct Data Structures	Proc. Scandinavian Workshop on Algorithm Theory (SWAT)	(PR)(CO)
31	2008	M. Olsen	The Computational Complexity of Link Building	Proc. International Conference on Computing and Combinatorics (COCOON)	(PR)
32	2008	M.A. Abam, M. de Berg and J. Gudmundsson	A Simple and Efficient Kinetic Spanner	Proc. Symposium on Computational Geometry (SoCG)	(PR)(CO)
33	2008	L. Arge, M.T. Goodrich, M. Nelson and N. Sitchinava	Fundamental Parallel Algorithms for Private-Cache Chip Multiprocessors	Proc. Symposium on Parallelism in Algorithms and Architectures (SPAA)	(PR)(CO)
34	2008	L. Arge, T. Moelhave and N. Zeh	Cache-Oblivious Red-Blue Line Segment Intersection	Proc. European Symposium on Algorithm (ESA)	(PR)(CO)
35	2008	P.K. Agarwal, L. Arge, T. Moelhave and B. Sadri	I/O-efficient Algorithms for Computing Contour Lines on a Terrain	Proc. Symposium on Computational Geometry (SoCG)	(PR)(CO)
36	2008	J. Feldman, S. Muthukrishnan, A. Sidiropoulos, C. Stein and Z. Svitkina	On Distributing Symmetric Streaming Computations	Proc. Symposium on Discrete Algorithms (SODA)	(PR)(CO)
37	2008	P. Indyk	Explicit Constructions for Compressed Sensing of Sparse Signals	Proc Symposium on Discrete Algorithms (SODA)	(PR)
38	2008	A. Andoni, P. Indyk and R. Krauthgamer	Earth Mover Distance over High-Dimensional Spaces	Proc. Symposium on Discrete Algorithms (SODA)	(PR)(CO)
39	2008	P. Indyk and A. McGregor	Declaring Independence via the Sketching of Sketches	Proc. Symposium on Discrete Algorithms (SODA)	(PR)(CO)

40	2008	K. Onak and A. Sidiropoulos	Circular Partitions with Applications to Visualization and Embeddings	Proc. Symposium on Computational Geometry (SoCG)	(PR)(CO)
41	2008	J. Matousek and A. Sidiropoulos	Inapproximability for metric embeddings into R^d	Proc. Symposium on Foundations of Computer Science (FOCS)	(PR)(CO)
42	2008	N. J. A. Harvey, J. Nelson and K. Onak	Sketching and Streaming Entropy via Approximation Theory	Proc. Symposium on Foundations of Computer Science (FOCS)	(PR)(CO)
43	2008	A. Andoni, D. Croitoru and M. Patrascu	Hardness of Nearest Neighbor under L-infinity	Proc. Symposium on Foundations of Computer Science (FOCS)	(PR)(CO)
44	2008	T. Chan, M. Patrascu and L. Roditty	Dynamic Connectivity: Connecting to Networks and Geometry	Proc. Symposium on Foundations of Computer Science (FOCS)	(PR)(CO)
45	2008	S. Mozes, K. Onak and Oren Weimann	Finding an Optimal Tree Searching Strategy in Linear Time	Proc. Symposium on Discrete Algorithms (SODA)	(PR)(CO)
46	2008	A. Chakrabarti, T.S. Jayram and M. Patrascu	Tight Lower Bounds for Selection in Randomly Ordered Streams	Proc. Symposium on Discrete Algorithms (SODA)	(PR)(CO)
47	2008	E. Demaine, T. Ito, Ni. J. A. Harvey, C. H. Papadimitriou, M. Sideri, R. Uehara and Yushi Uno	On the Complexity of Reconfiguration Problems	Proc. International Symposium on Algorithms and Computation (ISAAC)	(PR)(CO)
48	2008	E. Demaine, G. Aloupis, S. Collette, S. Langerman, V. Sacristan and S. Wuhrer	Reconfiguration of Cube-Style Modular Robots Using O(log n) Parallel Moves	Proc. International Symposium on Algorithms and Computation (ISAAC)	(PR)(CO)
49	2008	E. Demaine, M. Buadoiu, M. Hajiaghayi, A. Sidiropoulos and M. Zadimoghaddam	Ordinal Embedding: Approximation Algorithms and Dimensionality Reduction	Proc. International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX)	(PR)(CO)
50	2008	E. Demaine, T. G. Abbott, Z. Abel, D. Charlton, M. L. Demaine and S. D. Kominers	Hinged Dissections Exist	Proc. Symposium on Computational Geometry (SoCG)	(PR)(CO)

51	2008	E. R. Hansen, S. S. Rao and P. Tiedemann	Compressing Binary Decision Diagrams	European Conference on Artificial Intelligence (ECAI)	(PR)(CO)
52	2008	R. Berinde, P. Indyk and M. Ruzic	Practical Near-Optimal Sparse Recovery in the L1 Norm (invited paper)	Proc. Allerton Conference	(CO)
53	2008	R. Berinde, A. Gilbert, P. Indyk, H. Karloff and M. Strauss	Combining Geometry and Combinatorics: A Unified Approach to Sparse Signal Recovery (invited paper)	Proc. Allerton Conference	(CO)
54	2008	M.A. Abam, M. de Berg, and S-H. Poon	Fault-Tolerant Conflict-Free Coloring	Proc. Canadian Conference on Computational Geometry	(CO)
55	2009	R. Berinde, G. Cormode, P. Indyk and M. Strauss	Space-optimal Heavyhitters with Strong Error Bounds	Proc. Symposium on Principles of Database Systems (PODS)	(PR)(CO)
56	2009	V. Cevher, C. Hegde, P. Indyk and R. G. Baraniuk	Recovery of Clustered Sparse Signal from Compressive Measurements	Proc. International Conference on Sampling Theory and Applications (SAMPTA)	(PR)(CO)
57	2009	E. Demaine, G. Landau and O. Weimann	On Cartesian Trees and Range Minimum Queries	Proc. International Colloquium on Automata, Languages and Programming (ICALP)	(PR)(CO)
58	2009	D. Hermelin, G. M. Landau, S. Landau and O. Weimann	A Unified Algorithm for Accelerating Edit-Distance Computation via Text-Compression	Proc. International Symposium on Theoretical Aspects of Computer Science (STACS)	(PR)(CO)
59	2009	A. Kovacs, U. Meyer, G. Moruz and A. Negoescu	Online Paging for Flash Memory Devices	Proc. International Symposium on Algorithms and Computation (ISAAC)	(PR)
60	2009	G. Brodal, A. Jørgensen, G. Moruz and T. Mølhave	Counting in the Presence of Memory Faults	Proc. International Symposium on Algorithms and Computation (ISAAC)	(PR)
61	2009	D. Ajwani, A. Beckmann, R. Jacob, U. Meyer and G. Moruz	On Computational Models for Flash Memory Devices	Proc. Symposium on Experimental Algorithms (SEA)	(PR)(CO)

62	2009	U. Meyer and V. Osipov	Design and Implementation of a Practical I/O-efficient Shortest Paths Algorithm	Proc. Workshop on Algorithm Engineering and Experiments (ALENEX)	(PR)
63	20009	U. Meyer	Via Detours to I/O-Efficient Shortest Paths	Proc. Efficient Algorithms - Essays dedicated to Kurt Mehlhorn on the Occasion of his 60th birthday	
64	2009	D. Ajwani, R. Dementiev, U. Meyer and V. Osipov	Breadth First Search on Massive Graphs	Proc. Ninth DIMACS Implementation Challenge: The Shortest Path Problem	(PR)
65	2009	A. Beckmann, R. Dementiev and J. Singler	Building a Parallel Pipelined External Memory Algorithm Library	Proc. International Symposium on Parallel and Distributed Processing	(PR)
66	2009	G. S. Brodal and A. Jørgensen	Data Structures for Range Median Queries	Proc. International Symposium on Algorithms and Computation (ISAAC)	(PR)
67	2009	G. S. Brodal, R. Fagerberg, M. Greve and A. López-Ortiz	Online Sorted Range Reporting	Proc. International Symposium on Algorithms and Computation (ISAAC)	(PR)(CO)
68	2009	G. S. Brodal, A. Kaporis, S. Sioutas, K. Tsakalidis and K. Tsichlas	Dynamic 3-sided Planar Range Queries with Expected Doubly Logarithmic Time	Proc. International Symposium on Algorithms and Computation (ISAAC)	(PR)(CO)
69	2009	G. S. Brodal, A. Jørgensen and T. Mølhave	Fault Tolerant External Memory Algorithms	Proc. Algorithms and Data Structures Symposium (WADS)	(PR)
70	2009	A. Kaporis, A.N. Papadopoulos, S. Sioutas, K. Tsakalidis and K. Tsichlas	Efficient Processing of 3-Sided Range Queries with Probabilistic Guarantees	Proc. International Conference on Database Theory (ICDT)	(PR)(CO)
71	2009	M. Abam, M. de Berg, M. Farshi, J. Gudmundsson and M. Smid	Geometric Spanners for Weighted Point Sets	Proc. European Symposium on Algorithms (ESA)	(PR)(CO)
72	2009	M. Abam and M. de Berg	Kinetic Spanners in R^d	Proc. Symposium on Computational Geometry (SoCG)	(PR)(CO)

73	2009	M. Abam, P. Carmi, M. Farshi and M. Smid	On the Power of the Semi-Separated Pair Decomposition	Proc. Algorithms and Data Structures Symposium (WADS)	(PR)(CO)
74	2009	D. Ajwani	On P-complete Problems in Memory Hierarchy Models	Proc. Workshop on Massive Data Algorithmics (MASSIVE)	
75	2009	A. Farzan, R. Raman and S. Srinivasa Rao	Universal Succinct Representations of Trees?	Proc. International Colloquium on Automata, Languages and Programming (ICALP)	(PR)(CO)
76	2009	R. Pagh and S. Srinivasa Rao	Secondary Indexing in One Dimension: Beyond B-trees and Bitmap Indexes	Proc. Symposium on Principles of Database Systems (PODS)	(PR)(CO)
77	2009	R. Grossi, A. Orlandi, R. Raman and S. Srinivasa Rao	More Haste, Less Waste: Lowering the Redundancy in Fully Indexable Dictionaries	Proc. International Symposium on Theoretical Aspects of Computer Science (STACS)	(PR)(CO)
78	2009	J. E. Moeslund, P. K. Bøcher, J.-C. Svenning, T. Mølhave and L. Arge	Impacts of 21st Century Sea-level Rise on a Danish Major City – An Assessment Based on Fine-resolution Digital Topography and a New Flooding Algorithm	IOP Conference Series: Earth and Environmental Science 8	(PR)
79	2009	M. de Berg and P. Hachenberger	Rotated-Box Trees: A Lightweight c-Oriented Bounding-Volume Hierarchy	Proc. International Symposium on Experimental Algorithms (SEA)	(PR)(CO)
80	2009	P. Afshani, L. Arge and K. Dalgaard Larsen	Orthogonal Range Reporting in Three and Higher Dimensions	Proc Symposium on Foundations of Computer Science (FOCS)	(PR)
81	2009	P. Afshani, C. Hamilton and N. Zeh	A Unified Approach for Cache-Oblivious Range Reporting and Approximate Range Counting	Proc. Symposium on Computational Geometry (SoCG)	(PR)(CO)
82	2009	P. Afshani, C. Hamilton and N. Zeh	Cache-Oblivious Range Reporting With Optimal Queries Requires Superlinear Space	Proc. Symposium on Computational Geometry (SoCG)	(PR)(CO)

83	2009	P. Afshani, J. Barbay and T. Chan	Instance-optimal Geometric Algorithms	Proc Symposium on Foundations of Computer Science (FOCS)	(PR)(CO)
84	2009	L. Arge, M.T. Goodrich and N. Sitchinava	Parallel External Memory Model	Proc. Workshop on Theory and Many-Cores	
85	2009	L. Arge and M. Revsbæk	I/O-Efficient Contour Tree Simplification	Proc. International Symposium on Algorithms and Computation (ISAAC)	(PR)
86	2009	A. Andoni, P. Indyk, R. Krauthgamer and H.L. Nguyen	Approximate Line Nearest Neighbor in High Dimensions	Proc. Symposium on Discrete Algorithms (SODA)	(PR)(CO)
87	2009	A. Andoni, P. Indyk and R. Krauthgamer	Overcoming the L1 Non-embeddability Barrier: Algorithms for Product Metrics	Proc. Symposium on Discrete Algorithms (SODA)	(PR)(CO)
88	2009	R. Berinde and P. Indyk	Sequential Sparse Matching Pursuit	Proc. Allerton Conference	(PR)(CO)
89	2009	A. Andoni, K. Do Ba, P. Indyk and D. Woodruff	Efficient Sketches for Earth-Mover Distance, with Applications	Proc. Symposium on Foundations of Computer Science (FOCS)	(PR)(CO)
90	2009	A. Andoni, P. Indyk, K. Onak and R. Rubinfeld	External Sampling	Proc. International Colloquium on Automata, Languages and Programming (ICALP)	(PR)(CO)
91	2009	E. Demaine, M. Demaine, G. Konjevod and R. Lang	Folding a Better Checkerboard	Proc. International Symposium on Algorithms and Computation (ISAAC)	(PR)(CO)
92	2009	J. Cardinal, E. Demaine, M. Demaine, S. Imahori, S. Langerman and R. Uehara	Algorithmic Folding Complexity	Proc. International Symposium on Algorithms and Computation (ISAAC)	(PR)(CO)
93	2009	E. Demaine, M. Hajiaghayi, and D. Marx	Minimizing Movement: Fixed-Parameter Tractability	Proc. European Symposium on Algorithms (ESA)	(PR)(CO)
94	2009	B. Ballinger, D. Charlton, E. Demaine, M. Demaine, J. Iacono, C-H. Liu and S-H. Poon	Minimal Locked Trees	Proc. Algorithms and Data Structures Symposium (WADS)	(PR)(CO)

95	2009	E. Demaine, D. Kane and G. Price	A Pseudopolynomial algorithm for Alexandrov's Theorem	Proc. Algorithms and Data Structures Symposium (WADS)	(PR)(CO)
96	2009	T. Ito, M. Kaminski and E. Demaine	Reconfiguration of List Edge-Colorings in a Graph	Proc. Algorithms and Data Structures Symposium (WADS)	(PR)(CO)
97	2009	E. Demaine, M. Hajiaghayi and K. Kawarabayashi	Approximation Algorithms via Structural Results for Apex-Minor-Free Graphs	Proc. International Colloquium on Automata, Languages and Programming (ICALP)	(PR)(CO)
98	2009	E. Demaine, M. Hajiaghayi and P. Klein	Node-Weighted Steiner Tree and Group Steiner Tree in Planar Graphs	Proc. International Colloquium on Automata, Languages and Programming (ICALP)	(PR)(CO)
99	2009	E. Demaine, G. Borradaile and S. Tazari	Polynomial-Time Approximation Schemes for Subset-Connectivity Problems in Bounded-Genus Graphs	Proc. International Symposium on Theoretical Aspects of Computer Science (STACS)	(PR)(CO)
100	2009	E. Demaine, D. Harmon, J. Iacono, D. Kane and M. Patrascu	The Geometry of Binary Search Trees	Proc. Symposium on Discrete Algorithms (SODA)	(PR)(CO)
101	2009	E. Demaine, K. Kawarabayashi and M. Hajiaghayi	Additive Approximation Algorithms for List-Coloring Minor-Closed Class of Graphs	Proc. Symposium on Discrete Algorithms (SODA)	(PR)(CO)
102	2009	E. Demaine, M. Hajiaghayi, H. Mahini and M. Zadimoghaddam	The Price of Anarchy in Cooperative Network Creation Games	Proc. International Symposium on Theoretical Aspects of Computer Science (STACS)	(PR)(CO)
103	2009	J. Cardinal, E. Demaine, S. Fiorini, G. Joret, I. Newman and O. Weimann	The Stackelberg Minimum Spanning Tree Game on Planar and Bounded-Treewidth Graphs	Proc. Workshop on Internet and Network Economics (WINE)	(PR)(CO)
104	2009	J. McLurkin and E. Demaine	A Distributed Boundary Detection Algorithm for Multi-Robot Systems	Proc. International Conference on Intelligent Robots and Systems	(PR)(CO)
105	2009	G. Aloupis, N. Benbernou, M. Damian, E. Demaine, R. Flatland, J. Iacono and S. Wuhrer	Efficient Reconfiguration of Lattice-Based Modular Robots	Proc. European Conference on Mobile Robots	(PR)(CO)

106	2009	M. Ajtai, V. Feldman, A. Hassidim and J. Nelson	Sorting and Selection with Imprecise Comparisons	Proc. International Colloquium on Automata, Languages and Programming (ICALP)	(PR)(CO)
107	2009	R. Yuster and O. Weimann	Computing the Girth of a Planar Graph in O(n log n) time	Proc. International Colloquium on Automata, Languages and Programming (ICALP)	(PR)(CO)
108	2009	R. Backofen, G. Landau, M. Möhl, D. Tsur and O. Weimann	Fast RNA Structure Alignment for Crossing Input Structures	Proc. Symposium on Combinatorial Pattern Matching (CPM)	(PR)(CO)
109	2009	P. Klein, S. Mozes and O. Weimann	Shortest Paths in Directed Planar Graphs with Negative Lengths: A Linear-Space O(n log2 n)-Time Algorithm	Proc. Symposium on Discrete Algorithms (SODA)	(PR)(CO)

Journals

1	2007	G. S. Brodal, R. Fagerberg and G. Moruz	On the Adaptiveness of Quicksort	ACM Journal of Experimental Algorithmics, 12	(PR) (CO)
2	2008	D. Ajwani, T. Friedrich and U. Meyer	An O(n^{2.75}) Algorithm for Incremental Topological Ordering	ACM Transactions on Algorithms, 4(4)	(PR)
3	2008	M. Stissing, T. Mailund, C. N. S. Pedersen, G. S. Brodal and R. Fagerberg	Computing the All-Pairs Quartet Distance on a set of Evolutionary Trees	Journal of Bioinformatics and Computational Biology, 6(1)	(PR)(CO)
4	2008	L. Arge, M. de Berg, H. J. Haverkort and K. Yi	The Priority R-Tree: A Practically Efficient and Worst-Case Optimal R-Tree	ACM Transactions on Algorithms, 4(1)	(PR)(CO)
5	2009	M. Olsen	Nash Stability in Additively Separable Hedonic Games and Community Structures	Theory of Computing Systems, 45(4)	(PR)
6	2009	M. Abam, M. de Berg, M. Farshi and J. Gudmundsson	Region-Fault Tolerant Geometric Spanners	Discrete & Computational Geometry, 41(4)	(PR)(CO)
7	2009	M. Abam, M. de Berg and B. Speckmann	Kinetic kd-Trees and Longest-Side kd-Trees	SIAM Journal of Computing, 39(4)	(PR)(CO)
8	2009	L. Arge, V. Samoladas and K. Yi	Optimal External-Memory Planar Point Enclosure	Algorithmica, 54(3)	(PR)(CO)

9	2009	L. Arge, M. de Berg and H. Haverkort	Cache-Oblivious R-Trees	Algorithmica, 53(1)	(PR)(CO)
10	2009	H. Iben, J. O'Brien and E. Demaine	Refolding Planar Polygons	Discrete & Computational Geometry, 41(3)	(PR)(CO)
11	2009	E. Demaine, M. Hajiaghayi, H. Mahini, A. Sayedi-Roshkhar, S. Oveisgharan and M. Zadimoghaddam	Minimizing Movement	ACM Transactions on Algorithms, 5(3)	(PR)(CO)
12	2009	E. Demaine, M. Hajiaghayi and K. Kawarabayashi	Algorithmic Graph Minor Theory: Improved Grid Minor Bounds and Wagner's Contraction	Algorithmica, 54(2)	(PR)(CO)
13	2009	T. Abbott, M. Burr, T. Chan, E. Demaine, M. Demaine, J. Hugg, D. Kane, S. Langerman, J. Nelson, E. Rafalin, K. Seyboth and V. Yeung	Dynamic Ham-Sandwich Cuts in the Plane	Computational Geometry: Theory and Applications, 42(5)	(PR)(CO)

Thesis

1	2007	I. Brudaru	Heuristics for Average Diameter Approximation with External Memory Algorithms	MPI	MS Thesis
2	2007	G. Moruz	Hardware-Aware Algorithms and Data Structures	AU	PhD Thesis
3	2008	M. Patrascu	Lower Bound Techniques for Data Structures	MIT	PhD Thesis
4	2008	A. Sidiropoulos	Computational metric embeddings	MIT	PhD Thesis
5	2008	D. Ajwani	Traversing large graphs in realistic settings	MPI	PhD Thesis
6	2008	K. Do Ba	Testing closeness of distributions under the EMD metric	MIT	MS Thesis
7	2008	K. Lai	Complexity of Union-Split-Find Problems	MIT	MS Thesis
8	2008	J. M. Larsen og M. Nielsen	En undersøgelse af algoritmer til løsning af generalized movers problem i 3D	AU	MS Thesis
9	2008	C. Andersen	An optimal minimum spanning tree algorithm	AU	MS Thesis

10	2008	M. Revsbæk	I/O-efficient Algorithms for Batched Union-Find with Dynamic Set Properties and its Applications to Hydrological Conditioning	AU	MS Thesis
11	2008	A. H. Jensen	I/O-efficient Processing of LIDAR Data	AU	MS Thesis
12	2009	Martin Olsen	Link Building	AU	PhD Thesis
13	2009	Thomas Mølhave	Handling Massive Terrains and Unreliable Memory, AU	AU	PhD Thesis
14	2009	Henrik B. Kirk	Searching with Dynamic Optimality: In Theory and Practice	AU	MS Thesis
15	2009	Krzysztof Piatkowski	Implementering og udvikling af maksimum delsum algoritmer	AU	MS Thesis
16	2009	O. Weimann	Accelerating Dynamic Programming	MIT	PhD Thesis
17	2009	Volker Weichert	Radiation parameterization of the climate model COSMO/CLM in CUDA	FRA	MS Thesis
18	2009	R. Berinde	Advances in Sparse Signal Recovery Methods	MIT	PhD Thesis

Other

1	2008	E. Demaine, B. Gassend, J. O'Rourke, and G. T. Toussaint	All Polygons Flip Finitely ... Right?	In "Surveys on Discrete and Computational Geometry: Twenty Years Later", Contemporary Mathematics 453	(CO)
2	2008	A. Andoni and P. Indyk	Near-Optimal Hashing Algorithms for Approximate Nearest Neighbor in High Dimensions	Communications of the ACM, 51(1)	(CO)
3	2008	K. Mehlhorn and P. Sanders	Algorithms and Data Structures: The Basic Toolbox	Springer Verlag	(CO)

4	2009	D. Ajwani and U. Meyer	Design and Engineering of External Memory Traversal Algorithms for general	Algorithmic of Large and Complex Networks, Springer Verlag	(PR)
5	2009	L. Arge and N. Zeh	External-memory Algorithms and Data Structures	Algorithms and Theory of Computation Handbook, CRC Press	(PR)(CO)
6	2009	R. Hearn and E. Demaine	Games, Puzzles, and Computation	A.K. Peters	(CO)

Total number of publications in reported year	**Peer Rev.**	**Non-peer Rev.**
Number of journal articles	9	0
Number of conference series	52	3
Number of monographs	0	1
Number of book chapters	2	0
Others	0	7

List the 10 most prestigious channels of publication within the centres research area and prioritise in relation

Conferences

ACM Symposium on Theory of Computing (STOC)
IEEE Symposium on Foundations of Computer Science (FOCS)
ACM-SIAM Symposium on Discrete Algorithms (SODA)
Symposium on Computational Geometry (SoCG)
International Colloquium on Automata, Languages, and Programming (ICALP)
European Symposium on Algorithms (ESA)
ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)
International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX)/ International Workshop on Randomizationand Computation (RANDOM)
Scandinavian Workshop on Algorithm Theory (SWAT)/Workshop on Algorithms and Data Structures (WADS)
Workshop on Algorithm Engineering and Experiments (ALENEX)

Journals

Journal of the ACM
SIAM Journal on Computing
ACM Transactions on Algorithms
Discrete & Computational Geometry
Algorithmica
Journal of Computer and System Sciences
Computational Geometry: Theory and Applications
ACM Journal of Experimental Algorithmics
Theoretical Computer Science
Journal of Discrete Algorithms

State information about any bibliometric analysis the centre or an external part has made in accordance to standards within the centres field of research.

None

MADALGO List of Personel 2009 *01.01.-31.12.2009*

Personel		Hiring	Finansing (fraction of year) *)			Forign	For PhD and
Name	**Position**	**period in 2009 **)**	**Foundation**	**AU**	**Other finan-sing ***)**	**employee**	**Post Doc: Previous education**
Centerleder							
Lars Arge (AU)	Professor	all period	0,1	0,9			
Faculty							
Gerth S. Brodal (AU)	Associate Professor	all period		1			
Piotr Indyk (MIT)	Associate Professor	all period	0,1		0,1	x	
Eirk Demaine (MIT)	Associate Professor	all period	0,1		0,1	x	
Kurt Mehlhorn (MPI)	Professor	all period			0,1	x	
Ulrich Meyer (FRA)	Professor	all period			0,3	x	
Srinivas Rao (AU)	Post Doc	-31.01	0,1			x	MS, PhD
Mohammad Abam (AU)	Post Doc	all period	1			x	MS, PhD
Deepak Ajwani (AU)	Post Doc	all period	1			x	M.TECH, PhD
Peter Hachenberger (AU)	Post Doc	-01.06	0,5			x	MS, PhD
Peyman Afshani (AU)	Post Doc	01.02-	0,4		0,5	x	PhD
Nodari Sitchinava (AU)	Post Doc	01.08-	0,4			x	M.Eng., PhD
Gabriel Moruz (FRA)	Post Doc	all period			1	x	MS, PhD
Guests							
Jan van Leeuwen	Professor	13.05-25.06			0,1	x	
Norbert Zeh	Associate Professor	01.08-			0,4	x	
Martin Šmérek	PhD student	01.08-30.11			0,3	x	
Tehnical staff							
Adam E.N. Thomasen	Programmer	all period			1		
Thor Siiger Prentow	Programmer	all period	1				
Administrative staff							
Else Magård	Center manager	all period	1				
Ellen Lindstrøm	Accountant	all period		0,5			
Ph.d.-studerende							
Thomas Mølhave (AU)	PhD student	-31.08			0,7		BS + 1 år
Martin Olsen (AU)	PhD student	-30.07		0,3	0,4		Cand. Scient
Allan G. Jørgensen (AU)	PhD student	all period			1		BS + 1 år
Lasse Deleuran (AU)	PhD student	all period		0,7	0,3		BS + 1 år
Kostas Tsakalidis (AU)	PhD student	all period	0,4	0,6		x	BS + 1 år
Jesper Erenskjold Moeslund (AU)	PhD student	all period		1			BS + 1 år
Morten Revsbæk (AU)	PhD student	-31.10			0,8		Cand. Scient
Mark Greve (AU)	PhD student	all period		1			BS + 1 år
Pooya Davoodi (AU)	PhD student	all period	1			x	MS
Jacob Truelsen (AU)	PhD student	all period		1			BS + 1 år
Kasper D. Larsen (AU)	PhD student	all period		1			BS
Casper Kejlberg-Rasmussen (AU)	PhD student	01.02-		0,9			BS
Freek van Walderveen (AU)	PhD student	01.08-		0,4		x	MS
Sarah Zakarias (AU)	PhD student	01.08-		0,4			BS
Andreas Beckmann (MPI/FRA)	PhD student	all period	1			x	MS
Andrei Negoescu (MPI/FRA)	PhD student	all period			1	x	MS
Oren Weimann (MIT)	PhD student	-31.03.09			0,3	x	MS
Khan Do Ba (MIT)	PhD student	all period	0,6		0,4	x	BS
Jelani Nelsen (MIT)	PhD student	all period			1	x	MS
Eric Price (MIT)	PhD student	01.08-			0,4	x	BS
Morteza Zadimoghaddam (MIT)	PhD student	01.08-	0,2		0,2	x	BS

*) Approximation. Max one decimal.

**) More then three weeks.

***) Including no financing.