MADALGO
CENTER FOR MASSIVE DATA ALGORITHMICS

AARHUS UNIVERSITY

Danmarks Grundforskningsfond
Danish National Research Foundation

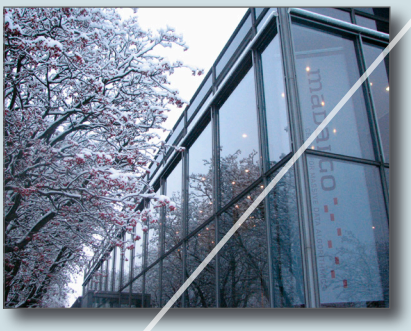

SELF EVALUATION REPORT

Center for Massive Data Algorithmics
2007 - 2010

The Danish National Research Foundation 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, traditional theory has used simplistic machine models that do not 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.

The goal of MADALGO is to be 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 center site at Aarhus University (AU) in Denmark (with Brodal and Arge as senior faculty) and at the sites of the other participants at Max Planck Institute for Informatics (MPI) and at Frankfurt University (FRA) in Germany (with Mehlhorn and Meyer as senior faculty) and at Massachusetts Institute of Technology (MIT) in the US (with Demaine and Indyk as senior faculty). The center also relies on international research collaboration, multidisciplinary and industry collaboration, and in general on a vibrant international environment at the center site. Another key is the focus on three related but also very different core research areas, namely *I/O-efficient*, *cache-oblivious* and *streaming algorithms*, as well as *algorithm engineering*. A more thorough discussion of the center motivation, objectives, core research areas and research strength of the senior faculty can be found in the center research plan.

This self evaluation report contains a summary of the center's activities and results in the first approximately 3 ½ years of the center period. As specified in the guidelines of the foundation, the report consists of three parts: *Part A* on research achievements, *Part B* on organization, leadership and collaboration and finally *Part C* on social impact of the center activities. The text of the individual parts includes answers to the specific questions put forward in the foundation guidelines.

Lars Arge
Center Director
September, 2010

A Research Achievements

The original center research plan discusses a number of main research challenges in each of the four core center research areas, and lists a number of concrete open geometric and graph problems. However, it also notes that the list of problems is non-exhaustive, and that the outcome of ongoing research of course should influence the exact directions taken in the center. In Section A1 below we first quickly recall the four core areas and highlight some of the main problems described in the original center research plan. In Section A2 we then discuss selected results in each of the areas. The original research plan also discusses that as the center matures, other interesting methodologies/directions for massive data processing than the four core areas should be considered. In the annual reports of the center we have indeed described several such new directions and research plan updates. In Section A3 we discuss these areas and some of the obtained results. Finally, we discuss center research education and larger events organized by the center in Sections A4 and A5. A summary and conclusions are given in Section A6.

A1 Original research plan

I/O-efficient algorithms are algorithms designed in a two-level 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. Even before the center was started the I/O-efficient algorithms area was quite developed. Not only had a large number of algorithms and algorithm design techniques been developed, but the immense practical importance of I/O-efficient algorithms had also been established through experimental work. Still many important problems remained open. The original research plan outlined a number of such problems, including a number of *fundamental geometric data structure* and *graph traversal problems* and some 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. The beauty of the model is that 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 a way of modeling a complicated (even unknown and/or changing) multi-level hierarchy using the simple two-level I/O-model. Unlike the I/O-efficient algorithms area, the cache-oblivious algorithms area is relatively new and even though efficient algorithms had been developed for a number of fundamental sorting and searching problems, as well as a few geometric and graph problems, at the center start, many even very fundamental problems remained open. The original research plan outlined a number of fundamental problems to be considered, mainly on *geometric data structure* and *batched problems (including terrain algorithms)*.

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 using significantly less space than the input data size, while processing each data element 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 accesses to disk blocks are much slower than sequential accesses. Moreover, in some applications data simply has to be processed sequentially as it is generated. The original research plan discussed how the streaming algorithms area has flourished as the discovery of several novel algorithmic techniques has enabled the enlargement of the class of problems with efficient streaming algorithms. Nevertheless, fundamental gaps remained in the understanding of what problems can be solved in the streaming model. The research plan thus outlined a number of fundamental problems to be considered, including investigation of the general *applicability of already developed techniques*, as well as of *fundamental geometric problems* and of *graph problems in variants of the streaming model*.

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 provides valuable input to work on theoretical models, and can lead to practical runtime

breakthroughs. The research plan outlined a number of algorithm engineering challenges in the I/O-efficient and cache-oblivious algorithms areas. In the *cache-oblivious* area focus was on *engineering simple algorithms for very fundamental problems*, whereas in the *I/O-efficient* algorithms area focus was 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*.

A2 Selected core research areas results

In general we have made steady progress on many of the problems discussed in the original research plan, as well as on several other problems. In this section we briefly discuss some of the obtained results. Note that due to space limitations, we will not be able to give a complete overview of all the obtained results. Appendix f contains a full list of center publications; the paper references below also refer to Appendix f, and the high-lighted (bold) papers are among the 10 selected publications listed in Appendix k.

I/O-efficient algorithms

In terms of **geometric data structures** we have obtained a number of interesting results in relation to the problems outlined in the research plan. We have e.g. developed an improved *point location* data structure, that is, a data structure for maintaining a dynamically changing subdivision of the plane such that the region containing a query point can be found I/O-efficiently [C29,JS6]. A number of structures have previously been developed for the problem, but using a number of new and existing techniques we managed to improve on the number of I/Os needed to update the subdivision. We have also developed new improved I/O-efficient structures for three and higher dimensional *orthogonal range reporting*, that is, the problem of storing a set of points such that the points in a query hyper-rectangle can be reported efficiently [**C80**,C125]. In three dimensions our new structures improve upon previously known structures and in higher dimensions they provide the first known non-trivial structures. In fact, using the techniques we developed, we were also able to design improved internal memory structures [**C80,C124**]. Furthermore, we also proved a lower bound that shows that our main *d*-dimensional data structure is space optimal. Recently, we have also obtained some improved structures (I/O-efficient as well as internal memory) for the 3-sided version of the problem with some randomness assumptions on the input points [C68,C120]. We have also studied sorted range reporting, where the input points have weights and where a query asks for the *k* points with smallest weight in the query range, either in an arbitrary or ascending weight order. We have provided an optimal structure for the unordered case and a query optimal structure for the ordered case in one dimension [CS11]. We have also proved a lower bound that shows that the space used by the latter structure is almost optimal.

In terms of **graph problems** we have also obtained a number of results along the lines of the original research plan. In particular, we have considered the problem of computing the *diameter* of a massive graph, which is a key challenge in complex network analysis. Since exact diameter computation is very costly, one is typically satisfied with an approximation of the diameter, often computed using a Breadth First Search (BFS) traversal of the graph (a traversal that visits all vertices at distance one from the source, then all vertices at distance two, and so on). However, even the best I/O-model BFS approaches suffer from poor performance. Hence, we investigated alternative approaches for approximating the diameter. Using controlled graph-contraction prior to a BFS or shortest-path traversal, we obtained improved algorithms both in terms of I/O-cost and approximation factor [C27]. We have also devised an approach to efficiently re-compute BFS after graph modifications (dynamic BFS) [C26]. The latter result implies an important separation in the context of **lower bounds**, namely that if BFS on sparse graphs cannot make efficient use of more than about the square root of the block size, then there is an intrinsic complexity gap between static and dynamic BFS in the I/O-model. This is not the case in internal memory. Still we have made less progress on I/O-efficient graph lower bounds than we would have liked.

As proposed, we have also worked on algorithms for **special graph classes**, for example on the intriguing problem of *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 of directed paths (related to the main memory and block size). For some classes of directed acyclic graphs we have shown that such a path cover can be obtained I/O-efficiently [CS3]. We have also obtained some results on the dynamic version of topological sorting [T5] (which in fact led to some internal memory improvements as well [J2,J18]). An overview of our (and other) achievements in I/O-efficient graph algorithms can be found in recent center surveys [C63,O4].

Finally, in terms of **terrain problems** we have also made progress on the problems outlined in the original research plan. We have developed an I/O-efficient algorithm for the problem of computing *contour lines* of a

terrain stored as a triangulation [C35]. This algorithm is the first I/O-optimal algorithm for the problem of computing all contours at a given height interval, as well as outputting these contours in clockwise order and with information about how they are nested. We have also developed an I/O-optimal algorithm for *simplifying the contour tree* of a terrain, which is a representation of the topological changes of a height h contour as h is increased [C85]. In previous work we showed how to compute the contour tree I/O-optimally [JA6].
In the area of terrain analysis, we have worked extensively on terrain water *flow modeling* problems, including the so-called terrain flooding problem. Traditional methods for terrain flow modeling do not account for water collecting in depressions in the terrain. Instead the terrain is typically flooded, that is, all depressions are removed, before computing how water flows over the terrain. Often this yields unrealistic flow patterns. More realistic predictions of water flow across the terrain can be obtained using "local terrain flooding" approaches, which remove "small" depressions before modeling flow. Using the techniques used in our contour tree simplification work, we managed to develop an I/O-efficient local flooding algorithm that can remove small depressions based on their height, area, or volume (or other geometric measures) [C85]. This result extends our earlier results on removing depressions based on their height [**C7**,JA6]. In very recent work we managed to account for depressions (even nested depressions) much more realistically by developing an I/O-efficient algorithm that computes exactly when a depression gets filled with water (assuming uniform rain over the terrain) based on its volume and upstream area [C127]. An overview of many of our terrain data processing results can be found in the PhD thesis of center student Mølhave [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 as discussed in the original research plan. However, our research has led to several important results in both directions.

In the area of cache-oblivious **data structures**, we have e.g. obtained several important results on cache-oblivious *dictionaries*, that is, data structures that allow for storing, updating and searching a collection of elements. Recently, we developed a cache-oblivious dictionary that uses the optimal number of cache misses for element insertions and deletions and for queries [**C119**]. 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. Other of our work explores tradeoffs between the insertion and search costs of cache-oblivious dictionaries [C11]. More precisely, we have shown how to speed up insertions significantly at the expense of a slight increase in the search cost.

A major research direction emphasized in our research plan is the development of techniques for constructing cache-oblivious **data structures for fundamental geometric problems**, including *range searching* problems in two and three dimensions. The results we have obtained in this area [C81,C82,JA7,JA8,CS10] make significant progress towards optimal data structures for such problems and further our understanding of the limitations of the cache-oblivious model. In [C81,JA8], we presented a general approach for obtaining cache-oblivious data structures for a class of range reporting problems that includes 2-d three-sided, 3-d dominance and 3-d halfspace range reporting. For 2-d three-sided range reporting, the obtained space and query bounds match those of previous data structures, but previously no comparable results were known for any problems in 3-d. The main challenge in obtaining reporting data structures for these problems turned out to be to obtain a data structure that can approximate the number of points in a query range. Just as previous 2-d three-sided range reporting data structures, our new range reporting data structures use more space than I/O-efficient data structures for the same problems. This raises the question whether this is unavoidable. In [C82,JA7], we gave an affirmative answer to this question by showing that any cache-oblivious data structure that achieves the optimal query bound for the above problems must use super-linear space. This result provides a strong separation between the I/O model and the cache-oblivious model but still leaves a substantial gap between the space lower bound we were able to prove and the size of existing cache-oblivious range reporting data structures (including those in [C81,JA8]). In collaboration with Norbert Zeh, who spent a sabbatical at the center in the fall of 2009, we recently made progress towards closing this gap. More specifically, we showed a new stronger space lower bound for optimal range reporting data structures, and we also substantially reduced the size of cache-oblivious data structures for 2-d three-sided and 3-d dominance range reporting [CS10].

We have also obtained important results in the area of **batched geometric problems**, another research direction outlined in the original research plan. We have e.g. developed a cache-oblivious algorithm for the *red-blue line segment intersection* problem, that is, the problem of finding all intersections between a set of non-intersecting

red segments and a set of non-intersecting blue segments in the plane [**C34**]. This problem has important applications in geographic information systems, where it arises in the context of map overlay operations. Our algorithm is optimal and, maybe even more importantly, the first cache-oblivious algorithm for an intersection problem involving non-axis-parallel objects. Finally, while our initial focus in the cache-oblivious model was on fundamental problems, we have recently started considering problems with more direct practical applications. For example, we have obtained results on the efficient *batched construction of* a so-called *quad-tree* on a set of points in the plane (and computing for each quad-tree leaf the set of points in all neighboring leaves) [M4]. This problem has applications in the computation of a grid terrain model from an irregular set of points.

Streaming algorithms

In the streaming algorithms area we have also made significant progress on the problems outlined in the original research plan. Our work can broadly be divided into two lines of research, namely the design of efficient algorithms for fundamental and widely applicable problems, and the identification of general frameworks and techniques that unify diverse results in streaming algorithms and beyond.

Along the first line of research, we have developed more efficient (and sometimes optimal) algorithms for a number of key problems. All of the problems can be represented in the following **geometric** manner: The stream of data elements is represented as a *count vector* with an entry for each possible element equal to the number of times the element occurs in the stream. The goal of an algorithm is then to compute some function of that vector under insertions (increment of an entry) and/or deletions (decrement of an entry). Since storing the count vector would require space equal to the total length of the stream, the goal is to store a "summary" of the vector instead, which is much shorter but nevertheless sufficient to estimate the desired quantity. The problems we have considered in this framework include the problem of *approximating the entropy* of a data stream. Entropy is a fundamental information-theoretic notion, which in the context of data streams is defined as the entropy of the distribution induced by the count vector. Several space-efficient algorithms have previously been designed for approximating the entropy, although not when both insertions and deletions are allowed. We managed to develop such an algorithm with a bound matching that for the insertions-only case [C42]. We also developed the first algorithms for the problem of space-efficiently estimating the degree of *independence* between two samples from joint distributions [C39]. For the well-studied problem of estimating the number of *distinct elements* in a stream, the number of non-zero entries in the count vector (its L0 norm), we developed the first optimal algorithm in a paper that received the Best Paper Award at the 2010 Symposium on Principles of Database Systems [C112]. For the similarly well-studied problem of estimating the L1 norm (under insertions and deletions), we have presented algorithms with space bounds matching the previously best known algorithms, but with substantially better running times [C115,C113]. Furthermore, some of our research on the distinct element problem and its generalizations turned out to be closely related to fundamental questions in pseudo-randomness [C114]. We have also considered the problem of finding *heavy hitters* in a stream, that is, the problem of finding the elements that occur frequently. The previously known algorithms fell into two classes, namely those with highly efficient space bound but sub-optimal accuracy guarantees, and those with higher space bounds but highly accurate estimations. We managed to obtain algorithms with the best of both worlds guarantees [C55]. Finally, as also proposed in the original research plan, we have considered the *bi-chromatic matching* problem, where we are given two sets of points (e.g., in the plane) of equal cardinality and the goal is to estimate the smallest cost of a one-to-one matching between the points in the first and the second set. This problem is also called the *Earth Mover Distance* problem. We provide the first know streaming algorithm with a constant factor approximation guarantee while preserving an efficient space bound [C89].
Finally, as also proposed in the original research plan, we have considered several problems in *multi-pass streaming models* where multiple passes over the data are allowed [C1,C16,T1].

For the second line of research, identification of **general frameworks and techniques**, our starting point was the "folklore" observation that *all* streaming algorithms supporting both insertions and deletions rely on an efficient *sketch* construction, that is, they naturally provide a way of summarizing data streams. We managed to show that this is not a coincidence [C36]. Specifically, we provided a general method for converting a streaming algorithm into a sketching scheme (for a wide class of functions that cover a large fraction of problems of interest). We also considered a fairly general class of *distance problems*, where the goal is to estimate the distance between two vectors, and provided general characterizations of distance functions that are solvable in the streaming model [C2].
One of the most powerful sketching methods is so-called *linear sketching*. For a count vector x, its *sketch* is equal to Ax, where A is a sketch matrix. Linear sketching is a technique used by many streaming algorithms,

including the algorithms for entropy, independence and L1 estimation mentioned above. It has also found many other applications e.g. in compressed sensing [C56] and efficient group testing [C111]. We focused on the recovery of a *sparse approximation* to x, that is, an approximation that has as few non-zero coefficients as possible while being as close to x as possible. This problem has been extensively investigated in recent years, notably in the area of compressed sensing. There are two major approaches to designing such sketches, namely, "combinatorial" and "geometric". The former typically leads to faster algorithms, while the latter results in shorter sketches. In our work, we showed that the two approaches are essentially different manifestations of the same phenomenon [**C53**]. In follow-up work, we further improve the recovery time to nearly linear in the dimension of x, while maintaining short sketches [C21,C52]. In fact, we showed that our algorithms achieve optimal sketch length [C110]. Our algorithms are the only known algorithms with both optimal sketch length and near-linear recovery. A recent center survey gives a broader overview of the area [J14].

Algorithm engineering

Following the original research plan, we have performed algorithm engineering work on our libraries for efficient implementation of I/O-model algorithms, on implementation of I/O-efficient graph algorithms, as well as on algorithms for massive terrain data. Algorithm engineering is also an integrated part of some of the work in the new directions areas discussed in Section A3. Unfortunately, we have not been successful in establishing a streaming code repository as proposed, just as we have yet to make real major progress on cache-oblivious algorithm engineering. However, we have experimented with our results on a cache-oblivious dictionary with insert/search tradeoff in the I/O setting, where we found that tremendous insertion speedups can be obtained compared to a normal (B-tree) structure [C11].

We have done **library work** on the TPIE and STXXL libraries for implementation of I/O-efficient algorithms. The TPIE work has included a general modernization of the library, as well as addition of features such as I/O-efficient priority queues. Ongoing work e.g. focuses on introducing pipelining (of a sequence of I/O-efficient sorting and scanning steps) in the library. The STXXL work includes internal memory parallelization and an augmentation of the pipelining feature to enable automatic task parallelization [C65].

In terms of **graph algorithms**, we have worked on the implementation of I/O-efficient algorithms for both Breadth First Search (BFS) and Single-Source Shortest Paths (SSSP). In connection with a (DIMACS) implementation challenge and based on STXXL, we have engineered some of our previous theoretical BFS approaches to obtain an algorithm that in some cases runs up to three orders of magnitude faster than the previous best solutions [T5,C64]. We have also engineered a simplified SSSP algorithm, and reported on experimental results on undirected sparse graphs under the assumptions that the ratio between the largest and the smallest edge weight is reasonably bounded, and that the main memory is large enough to keep one bit per vertex [C62]. While our implementation only guarantees average-case efficiency, i.e., assuming randomly chosen edge-weights, its performance on real-world instances with non-random edge weights is even better than on synthetic inputs with random weights. On a real-world 24 million node street graph, our implementation was over 40 times faster than previous implementations. Our graph traversal implementations, our theoretical work on approximate BFS [C27], and our alterative heuristic [T1] are important ingredients in ongoing work towards implementations of I/O-efficient diameter approximation methods [M2]. An overview of much of our work on implementing I/O-efficient graph algorithms can be found in a recent center survey [O4].

We have also worked extensively on engineering I/O-efficient algorithms for massive **terrain data processing**. In fact, we have reimplemented and significantly extended the capabilities of our TerraFlow terrain processing software package based on TPIE. Over the center period, a number of new (limited) releases have been made of the resulting TerraStream software package to a number of direct research and industry collaborators, as well as on request to more than 30 other institutions and companies. Many of the changes in the new releases – new smaller feature and bug fixes – have been made as a result of feedback from users. Some of the larger new features include implementations of I/O-efficient algorithms for constructing terrain models in the standard so-called grid and TIN (triangulated irregular network) formats from massive terrain point samples [**C7**,C126]. TerraFlow only supported terrains stored as a regular grid of height values. Some of the grid model work was performed in collaboration with Assistant Professor Andrew Danner from Swarthmore College in the US while he was visiting the center for a month in 2008, and some of the work includes a new method for estimating the quality of the computed model [M3,M4]. We have also worked on the problem of removing outlier points from "raw" terrain data before constructing terrain models, mainly in connection with massive sonar point datasets, where outliers appear e.g. as a result of scans of fish, multiple sonar reflections or scanner self reflections [CS5].

This work is performed in collaboration with industry experts in EIVA A/S, Statoil and a few other companies. Other new TerraStream features include implementation of erosion modeling algorithms (so-called LS-factor computation) and a tool for predicting where flooding is likely to occur in case of a rise in ocean level. Following our theoretical results on contour line computation and on removing small depressions in a terrain based on any of a number of local geometric measures (the flooding problem) described above, we have also worked on engineering practical algorithms for these problems [C85]. Our solution to the flooding problem has proven particularly efficient and interesting, not only in connection with realistic flow modeling but also e.g. in connection with simplification of contour lines. An overview of some of the uses of our new flooding algorithms can be found in the MS thesis of center PhD student Revsbæk [T10]. Overall, TerraStream now consists of a number of software modules that allow for a whole terrain data processing "pipeline" from removing of outliers from raw data, through terrain model construction and terrain simplification (using flooding), to more complicated water flow analysis tools. To our knowledge, TerraStream is the only software package that is capable of performing these computations on massive countrywide terrain datasets, and it (especially the parts pertaining to water flow and flood prediction) has thus attracted a lot of attention from industry and practitioners [**C7**,C126]. Using the software and very detailed terrain data from industry collaborator COWI A/S, interdisciplinary center PhD student Moeslund (who is a Biology student co-advised by MADALGO center director Arge) has e.g. conducted a study of the impact of a sea-level rise on Aarhus [C78], which received quite a lot of attention from the local media. Recently he has worked on the relationship between vegetation and elevation in salt meadows across Denmark and the implications for sea-level rise impacts [JS4]. Details about his work can also be found in his recent MS thesis [T20].

Finally, it is important to note that while our theoretical work on I/O-efficient algorithms obviously inspired a lot of the above algorithm engineering work, our engineering work also continues to inspire theoretical work. Most recently, it has e.g. inspired some interesting results on internal memory line and polygon simplification [CS4,M9]. Some of these results are in collaboration with visiting PhD student Shervin Daneshpajouh.

A3 New directions areas and results

The exact research directions of a center like MADALGO are of course highly influenced by the research progress being made, as well as by the interests of center researchers. Already in the original research plan it was discussed that as the center matures, we will consider other interesting or even more realistic methodologies/directions for massive data processing than the core research areas; one of the external proposal reviewers also commented that we should maybe consider methodologies that incorporate parallel computation and that we could consider so-called succinct data structures. During the center period, we have in fact explored a number of new directions and considered a number of new methodologies/models, such as *parallel private-cache*, *fault tolerance* and *flash memory models*, as well as *succinct data structures*. We have also had success with work that cuts across individual models. Much of the focus of our work has been on the theoretical possibilities and limitations of different models and thus technologies, while the overall goal of course is to develop better and more realistic theoretical models for efficient algorithm design. Another crosscutting "theme" that has emerged in the center's work is *data structures*, that is, efficient ways of storing (massive) data using small space, such that queries on the data can be answered efficiently. Much of the work in the core areas described above, as well as in the new directions areas described below, is on data structuring problems. Below we further discuss the new directions and themes that have evolved during the center period, and highlight some of the obtained results.

Parallel private-cache model

Parallelism is an increasingly popular way to handle massive data and although not a main focus area of the center, we are of course considering parallel algorithm issues whenever relevant. As mentioned above, we have e.g. worked on parallel algorithms in the context of algorithm engineering. With the aim of e.g. better understanding the relationship between parallel and I/O-efficient algorithms, we have also done some work on the intriguing connections between provably efficient parallel algorithms and streaming- and I/O-model solutions for the same problems [C74,C63].

The bulk of our parallel algorithm work has been on *multi-core* algorithms. Chip manufacturers are increasingly producing chips with several CPUs (or *cores*) on a single chip; current architectures have 2, 4 or 8 cores, but it is predicted that this number will grow dramatically in the not too distant future. Thus there is a need for parallel algorithms that can utilize the many cores and for theoretical work to characterize the opportunities and limitations of multi-core architectures. Such work might also eventually influence future multi-core architec-

tures. This need has also been recognized by the broader algorithms community; for example, multi-core algorithms were one of the themes of the (biannual) 2010 Dagstuhl seminar (invitation only seminars held at the Dagstuhl castle in Germany) on data structures (co-organized by center director Arge).
Most of the algorithms described in the vast literature on parallel algorithms are developed in models that do not adequately describe the new multi-core architectures, especially their cache-architectures. In the center annual reports we have discussed challenges and work in a more adequate *parallel private-cache model* that 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 are simultaneously transferred between the external memory and each of the main memories. This model is one of a few theoretical models that try to capture modern multi-core architectures. Our work on parallel private-cache algorithms was initiated when UC Irvine Professor Michael Goodrich visited the center as a Fulbright scholar in 2007 along with his PhD student Nodari Sitchinava. Sitchinava visited again in the summer of 2008 and became a center Post Doc in 2009. Initially we obtained efficient algorithms for several *fundamental problems* such as selection and sorting [C33]. By adapting several I/O-algorithm techniques for solving graph problems efficiently to the parallel setting, we have also managed to develop efficient algorithms for several *graph problems* in the model [C123]. Recently, we have investigated which I/O-efficient design techniques can be used to obtain efficient 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. However, we recently showed how to combine the sequential I/O-efficient solutions with techniques from our previous work on parallel private-cache model algorithms to obtain efficient algorithms for orthogonal line segment intersection and other geometric problems [**C129**].

Models of flash memory
Flash memory devices are becoming increasingly large and cheap. They may eventually replace (or at least supplement) disks as the external storage in mobile computing. Since flash memory appears to have very different characteristics than both internal memory and disks, it is important to try to characterize the new devices in order to develop theoretical models of flash memory and thus to understand the possibilities and limitations of these devises, and obviously also to develop efficient algorithms for them. These theoretical considerations may ultimately influence the exact architecture of future flash devices. As in the case of multi-core algorithms, the importance of theoretical work in the area of flash memory has been recognized by the algorithms community, and flash memory was one of the themes of the 2008 Dagstuhl seminar on data structures (also co-organized by center director Arge).
In the annual reports we have discussed our work on models of flash memory and its connection to mainly our work on I/O-efficient algorithms. Initially, we worked on characterizing the performance of various flash-based storage devices [C25]. Besides analyzing an expected huge difference in read and write speed, we also e.g. analyzed the effects of read/write patterns on performance. It turned out that despite the similarities between flash memory and internal memory (fast random reads) and between flash disks and normal disks (data block movement), algorithms designed in internal memory models or the I/O-model do not necessarily realize the full potential of flash memory devices. In fact, just replacing a hard disk by flash may result in degenerating performance. However, we also showed that a careful combination of internal memory and I/O-model algorithmic ideas can yield nice speedups. Subsequently, we used our experimental work to propose two theoretical models of flash memory: The *general flash model* and the *unit-cost flash model* [C61]. The general flash model is similar to the I/O-model, except that read and write block sizes are different and incur different costs. The unit-cost flash model augments the general flash model with the assumption that the throughput provided by sequential reads and writes is equal. We also showed that a large body of existing merging-based I/O-model algorithms can be easily adapted to the unit-cost flash model [C61], and we are currently investigating if efficient algorithms can also be obtained for problems that are solved using distribution in the I/O model. Finally, we have proposed a variation of standard *online paging* in two-level memory systems geared towards flash memory devices [C59], as well as considered algorithms that not only minimize cache-misses but also the time spent on deciding what block (page) to evict when a new block is loaded [CS7].

The use of I/O-efficient algorithms techniques to develop algorithms that are efficient in the various models of flash memory, as well as to develop parallel private-cache model algorithms, are two examples of some of our work that cuts across different models of computation (although it might not be particularly surprising since both the flash and parallel private-cache models are inspired by the I/O-model). Recently, we also had success

with an algorithm engineering effort that combines our work in I/O-efficient algorithms, parallel processing and flash memory, namely *energy efficient sorting*. Using a low-power processor, solid flash disks, and I/O-efficient algorithm techniques (as well as techniques from several other areas), we managed to beat the current energy efficient sorting world records in several categories of the so-called JouleSort benchmark [**C116**]. For example, in the category of sorting 1 Terabytes of data with random keys, we improved the previously best record by more than a factor of five. Our record got worldwide media coverage.

Fault tolerance models

Modern memory is not always fully reliable; sometimes the content of a memory location may be temporarily or permanently corrupted. This may depend on manufacturing defects, power failures, or environmental conditions such as cosmic radiation. Furthermore, error rates are expected to increase as memory is getting smaller and more complex, and working at lower voltage and higher frequencies. Thus memory errors may become a serious problem when processing massive data, since it normally means performing a large number of memory accesses on many memory devices and over a long period of time. Therefore we have, as discussed in the annual center reports, also considered how to handle memory errors algorithmically in massive data computations.

Traditionally, the handling of errors has been addressed in various ways. At the hardware level they are often tackled using various error detection mechanisms (such as redundancy, parity checking or Hamming codes). However, such mechanisms often involve non-negligible penalties with respect to performance, size, and cost. Thus memories implementing these mechanisms are rarely found in ordinary workstations. In the algorithmic community dealing with unreliable information has been addressed in a variety of different settings. Very recently, one particularly interesting model called the *faulty-memory RAM model* has been proposed for modeling and handling memory errors. This model extends the classical internal memory RAM model in the sense that memory cells may get corrupted at any place and at any time (even simultaneously) during the execution of an algorithm; corrupted and uncorrupted cells cannot be distinguished. We have obtained a number of important results in this model. For example, we have developed the first optimal *dynamic data structures* for the model, namely dynamic dictionaries [C6] and priority queues [C5], and studied the trade-off between the accuracy of maintaining a *collection of binary counters* and the increment time of a counter [C60]. In the latter work 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. Recently, we have also considered fault-tolerant *kd-trees* for use in clustering algorithms from a theoretical as well as an algorithm engineering point of view [CS1].

The early work on fault tolerant algorithms and data structures does not scale beyond internal memory. Thus, just like our parallel private-cache and flash model work has involved the use of I/O-efficient algorithm techniques, we have recently worked on combining fault tolerance and I/O-efficiency. More precisely, we have proposed a faulty-memory version of the I/O-model and developed optimal *sorting* algorithms, as well as dictionary and priority queue *data structures*, in the model [C69]. For the case of dictionaries the optimality of our structures was proved by a new lower bound that is stronger than what can be derived from either existing I/O-model lower bounds or faulty-memory RAM lower bounds. Our work has revealed some interesting dependencies, for example that the optimal bounds depend on whether faults can happen both in internal and external memory, and on if the algorithms are deterministic or randomized. The PhD thesis of center student Jørgensen gives an overview of our faulty-memory results [T19].

Succinct data structures

When dealing with massive data it can be important to use really space-efficient data structures, for example in order to be able to transport them efficiently or store them in the main memory of a relatively small device. However, it is of course also important that the reduced space does not come at a cost of e.g. much higher query complexity. 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. Following the recommendation of one of the center proposal reviewers, we have considered the area of succinct data structures.

The fundamental and well-studied problem of *encoding a bit-vector* space-efficiently forms a basic building-block in many succinct data structures. The problem consists of representing a static bit-vector such that a rank operation (counting the number of ones up to a given position) and a select operation (finding the position of the *i*th one in the vector) can be supported efficiently. We have studied this problem in two different models and obtained several new results. In the indexing model, where the bit-vector is explicitly stored together with an indexing data structure, we have generalized existing lower bounds and obtained optimal structures [C30]. In

the encoding model, where the bit-vector is not explicitly stored, we improved the redundancy of the data structure (i.e. the space needed in addition to the information theoretic lower bound) by a logarithmic factor for dense vectors [C19,C30]. We also generalized our structures to handle strings over larger alphabets, i.e. to perform rank and select queries with respect to specific symbols in the vector. Using a non-trivial recursive procedure, center PhD student Patrascu further improved this space redundancy by a poly-logarithmic factor [**C23**]. The result solved a long-standing open problem in succinct data structures, and Patrascu received the best student paper award at the 2008 Symposium on Foundations of Computer Science (one of the top theoretical computer science conferences) for the paper describing the results.
The bit-vector problem is closely related to the so-called predecessor problem (i.e. to find the position of the first 1-bit to the left of a given position). Recently, we further tightened the connections between the two problems and gave an improved structure for the so-called fully indexable dictionary problem, that is, the problem of storing a bit-vector supporting rank and select operations on both 0-bits and 1-bits [C77]. We have also considered succinct representations of Binary Decision Diagrams (BDD) [C51], used in formal systems verification, as well as improved the representations of ordered trees [C75], dynamic sets [CS2] and multi-sets [JS3]. An experimental study of our BDD structures showed a space improvement of up to a factor of 4 over previous structures.

Finally, like in our work on fault tolerant algorithms, we have also worked on combining I/O-efficient and succinct data structures. More precisely, we developed the first static space- and query-optimal external memory search tree structure supporting range queries [C76]. Earlier known structures (namely B-trees and compressed bitmap indexes) did not achieve both space and query-time optimality simultaneously. For the dynamic version of the problem we presented algorithms allowing a tradeoff between the update time and the space usage.

Data structures
As discussed in the annual reports, an emphasis on data structures is a crosscutting "theme" that has emerged in the center. Not only is much of the center work in the various models described above on data structuring problems (e.g. [C5,C6,C11,C19,C22,**C23**,C28,C29,C44,C45,C51,C60,C68,C69,C75,C76,C77,**C80**,C81,C82, **C119**,C120,C125,J2,J4,JA3,JA7,JA8,JS6,CS1,CS2,CS10,CS11,M6]), but the broad data structure expertise in the center also represents an opportunity to consider even classical open data structure problems.

Our recent work on the so-called dynamic optimality conjecture and on higher-dimensional orthogonal range reporting are two examples of center work on classical data structure problems. The *dynamic optimality conjecture*, that is, the conjecture that splay trees (or any other online search tree data structure) perform within a constant factor of an optimal offline search tree data structure for all possible access sequences, is more than 25 years old. We recently gave a new perspective on the conjecture by studying the combinatorial properties of search trees as a geometric problem [C100]. This work was performed in collaboration with Associate Professor John Iacono from Polytechnic Institute of NYU, who visited the center for an extended period of time in the summer of 2008 and again in 2010. The *higher-dimensional orthogonal range reporting* problem is also a classical and longstanding open problem. 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. Using I/O-efficient algorithms techniques, we managed to develop a space and query optimal three-dimensional structure and showed that the query time indeed has to increase with dimension [**C80,C124**]. The work opens up some intriguing new problems. For example, the optimal two- and three-dimensional structures answer 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. We have also considered various other one- and two-dimensional internal memory range searching problems. In one dimension we have presented space optimal data structures for returning the elements in a subrange of an array in sorted order [C69], data structures for returning the median element in a subarray [C66,JA3,CS9], and for finding the most frequent element (mode) in any subarray [C117]. It is worth noting that the latter result 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,C128].

A4 Center research education

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. Currently, the center houses 18 PhD students (11 at AU, 3 internationals) and 6 Post Docs (5 at AU, all internationals). Furthermore, 9 center PhD

students (5 at AU) have already graduated. A list of center Post Docs and PhD students can be found in Appendix b and c, respectively. Currently, 5 MS students (3 at AU) are also affiliated with the center and 13 master degrees (8 at AU) have already been obtained at the center.

The center's focus on research education includes exchange of Post Doc and PhD students, and almost all MIT and MPI/FRA Post Docs and PhD students have visited AU several times. AU PhD students also stay around 6 months abroad during their PhD study. So far these stays have included stays at TU Berlin (Nilsson), University of Rome (Moruz), ICT Sidney (Olsen), AT&T research (Mølhave), Duke University (Mølhave and Deleuran), MIT (Jørgensen), University of Waterloo (Tsakalidis), Leicester and Carleton (Davoodi). Currently planned stays include stays at MIT (Tsakalidis), ETH-WSL (Moeslund) and Princeton (Larsen). Regular center research seminars and workshops are obviously also an important part of research training, and during almost the entire center period the center has had weekly internal research seminars at AU with research talks given by center employees and visitors. The PhD students have also run a student (and Post Doc) only seminar, just as the center has hosted a number or "open problems sessions" and smaller workshops (e.g. around the TerraStream software). Organization of advanced courses and schools is also an important part of research education. Apart from a number of regular graduate (and undergraduate) level classes, the center Post Docs and visitors have taught 4 specialized and intensive 5 ECTS PhD classes at AU, and as discussed in Section A6 below the center has also organized three large international summer schools (center researchers have taught at several other international summer schools). A list of center educational activities is included in Appendix i. Regular research visits from and collaboration with non-center researchers are important ingredients in the center's efforts to create a world-leading and international environment, just as participation in international conferences and workshops outside the center are important for research training. Center collaborations are discussed further in Section B1, and conference participation with contribution is listed in Appendix h. Finally, the center emphasizes initiatives designed to create a sense of center community. This includes a yearly two day fall retreat, monthly center lunches at AU, as well as a number of social events (Christmas dinner, weekly breakfast, weekly soccer, go-carting trips, etc) and various "merchandise" such as MADALGO coffee mugs, t-shirts and caps (also used for publicity purposes at the larger center organized events). A "research pipeline board" is also maintained at AU, where one can follow the progress on AU center research papers.

A5 Larger center research events

During the center period, the center has organized a number of larger events, starting with a four day international summer school in August 2007 on *data stream algorithms*. The lecture program was coordinated by senior center researcher Indyk, who also was one of the lecturers at the school. The school had around 70 participants (mostly PhD students) representing 21 nationalities. Immediately following the school the center organized a center inauguration event, which e.g. featured scientific talks by highly-recognized (non-center) international researchers in the core center research areas. In August 2008 the center organized another four day international summer school, this time on *cache-oblivious algorithms*. The lecture program was coordinated by senior researchers Brodal and Demaine, who were also among the lecturers at the school. The school had around 60 participants (mostly PhD students) representing 21 nationalities. In June 2009 the center organized the 25th Annual Symposium on Computational Geometry (SoCG) at AU. 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 the center organized a celebration of the 25th anniversary of the symposium. The day after the symposium the center organized the first Workshop on Massive Data Algorithmics (MASSIVE). 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 workshop, the center organized a second workshop in connection with the 2010 SoCG in Snowbird, Utah, USA. This year the workshop was attended by 30 researchers (the 2010 SoCG attendance was also down relative to 2009) and featured 14 contributed talks. The hope is to eventually make MASSIVE a full-fledged conference, probably co-located with one of the broader algorithms conferences. Finally, in August 2010 the center organized yet another international summer school on *geometric data structures*. The lecture program was coordinated by John Iacono in collaboration with senior center researchers Arge and Brodal. Iacono was also one of the lecturers at the school, which had around 50 participants (mostly PhD students) from 26 different institutions in 14 countries. Judging from the evaluations conducted at each of the summer schools and several of the other events, all of the large center events were very successful. More information about the events can be found at www.madalgo.au.dk.

A6 Summary and conclusions

Overall, we believe the center has been quite successful so far. As outlined above, we have made great progress in the four core center research areas. In fact, we have made progress on almost all of the problems outlined in the original research plan; we have developed several new I/O-efficient data structures, obtained new I/O-efficient results on various graph traversal problems (including for special graph classes), and made major progress on I/O-efficient algorithms for various massive terrain data problems; we have developed new and improved cache-oblivious data structures, including for several geometric problems, and obtained cache-oblivious algorithms for several batched geometric problems; we have obtained major progress on the general applicability of general streaming algorithm frameworks and techniques, developed new and improved algorithms for numerous streaming problems, and also considered algorithms in multi-pass and semi-streaming models; we have further developed our I/O-efficient software libraries, engineered I/O and practically efficient algorithms for several fundamental problems, and made major progress on our software for efficient processing of massive terrain data. Much of our work has been performed with a large number of other researchers (discussed further in Section B1), including in multidisciplinary collaborations, and with a number of industry collaborators (discussed further in Section C1). We believe the only major areas of the original research plan where have not made real progress is on establishing a central repository for streaming algorithm implementations, just like we have made less progress on engineering cache-oblivious algorithms and on I/O-efficient graph algorithm lower bounds than we would have liked. Compared to the goal outlined in the original research plan, we have also not yet established a really strong presence in the streaming algorithms area at AU, although progress has been made through senior researcher Indyk's visits, the organization of a streaming algorithms summer school, as well as the hiring of AU Post Doc Abam (in 2008 and 2009) with some streaming algorithm expertise. However, newly hired Post Doc Zhang adds significant strength in the area at AU.

The original research plan has been continuously adjusted to e.g. include further methodologies for massive data processing than the four core areas. Thus as described, we have obtained significant results on algorithms for multi-core (parallel private-cache) architectures and for flash memory, as well as algorithms that are tolerant to faults, and on space efficient (succinct) data structures and data structures in general. As mentioned, we have also had success with work that cuts across individual models: We have e.g. considered succinct data structures in the I/O-model, the combination of fault tolerance and I/O-efficiency, the use of I/O-algorithm techniques in the design of parallel private-cache and flash model algorithms, as well as in the design of internal memory data structures, just as we have considered the connections between parallel, streaming and I/O-models solutions for the same problem (or classes of problems). The work on these new models and their combination fits well with the overall motivation of the center, namely to consider algorithms in theoretical models that more adequately model modern applications. Ultimately, our dual focus on investigating the basic theoretical possibilities and limitations as well as on developing practically efficient algorithms, might lead not only to more realistic models of existing hardware (like e.g. in the case of flash memory) and practical runtime breakthroughs (like e.g. in the case of massive terrain data processing), but also influence future hardware architectures.

In terms of center activities and concrete milestones we have also followed the original research plan to a large extent. As planned, we have e.g. made several software releases, arranged several summer schools and smaller workshops, and initiated the MASSIVE workshop. We decided not to invest in dedicated video conference facilities between the sites (after a cost-benefit analysis), just as we decided to focus on maintaining the center's web-portal rather than to establish a newsletter and a technical report series. We have also found that we could attract research visitors to AU without a formalized visiting fellow and summer student program. Finally, although we have produced some educational material, including lecture notes in connection with the center summer schools, we would like to intensify this work.

In terms of our goal of training the next generation of researchers in a world-leading and international environment, we also believe our efforts have been relatively successful. A large number of Post Doc and PhD students is and has been associated with the center, and the center has had a focus on exchange, collaboration and training of PhD students. We also believe our Post Doc recruiting efforts have been quite successful, and we are experiencing high interest when we advertize Post Doc positions. Note also that three PhD students at one center site have gone on to Post Doc positions at another center site (Moruz from AU to FRA, and Hachenberger and Ajwani from MPI to AU), and that one center AU Post Doc (Sitchinava) visited the center twice as a PhD student. A good deal of our Post Docs and PhD students have also gone on to obtain good research positions (refer to Appendix b and c). In terms of PhD student recruiting, we believe we have recruited some very strong students (including several students in the recently introduced AU "honors" program).

However, we would have liked to have more international students at AU. While all our Post Docs are internationals, only 5 out of the 17 center PhD students at AU are internationals. We believe there is e.g. an opportunity to recruit more highly skilled Iranian students – now that we have already recruited two Iranian Post Docs and a PhD student – as well as Eastern European students. Although difficult, we would also like to recruit female PhD students. It should be noted that the number of center PhD students is actually higher than originally planned, and that as planned most of the students were admitted early in the center period.

Overall, we believe that the accomplishments of the center compare very favorably to other high standing research groups in the field internationally. In fact, we believe that the center is the top research institution world-wide in I/O-efficient algorithms and possibly also in cache-oblivious algorithms. The center is also among the top research centers in algorithm engineering and streaming algorithms. Thus we believe that the center has made great progress towards the goal of being a world-leading center in algorithms for handling massive data. We believe the structure of the center that brings together a strong set of senior researchers from AU, MIT, MPI and FRA with complementary strengths has played a central role in the success of the center. The collaboration between the sites has lead to several concrete research results, but the strength of the team also obviously played a major role in heightening the international visibility of the center (and thus e.g. for recruiting), just as it has played an important role in establishing a world-leading international environment at AU, for organizing events like the three summer schools, and so on. One indicator of the international standing of the center is that during the center period, center researchers have given more than 30 talks at invitation-only international research workshops, given more than 70 invited presentations at international research conferences and research (university or industry) seminars, as well as given more than 70 regular presentations at international research conferences (refer to Appendix h); according to Google Scholar, center publications have already been cited 1195 times (refer to Appendix f). Another indicator of the international standing is the many awards and acknowledgments received by center researchers (listed in Appendix h). These include two honorary degrees, two best paper awards and two associate editor appointments. Center Director Arge was also elected member of the Royal Danish Academy of Science and Letters and elected a Distinguished Scientist by the Association of Computing Machinery. He also recently received an Elite Researcher Award from the Danish Minister of Research.

B Organization, Leadership and Collaboration

B1 Center research collaboration

By design the center is highly collaborative and one of the main goals of the center is to maintain a vibrant, world-class and international environment at AU. Thus emphasis has e.g. been on hosting international visitors (faculty as well as PhD students) at AU. The senior MIT, MPI and FRA center researchers have visited AU several times, including for longer periods of time, especially in connection with summer school organization. As mentioned, almost all of the MIT and MPI/FRA PhD students and Post Docs have also visited AU several times. As indicated in Appendix d, several non-center faculty and PhD students have also visited AU for longer periods of time (more than 3 weeks). Just as important as long term visitors, the center has also hosted many shorter term visitors (several for more than one period of time). These visitors are also listed in Appendix d. In general, center researchers have extensive research collaboration with other computer scientists, as can be seen from the list of collaborators (most with joint publications) found in Appendix h. As discussed in Appendix f, center researchers have published with 192 non-center researchers in the center period, and only 62 of the 235 center publications are authored solely by center researchers.

Through collaborative efforts, the center has also tried to be a catalyst for multidisciplinary collaboration (as well as industry collaboration, as discussed further in Section C1). Thus the center has collaborated extensively with environmental and agricultural researchers at the Institute of Biology, the Faculty of Agricultural Sciences and the National Environmental Research Institute at AU, as well as at Duke University and NC State University, on issues in connection with massive terrain data. As discussed in Section A2, this collaboration has been centered around the TerraStream software package and also funded through several other sources such as a NABIIT grant from the Danish strategic research council, a grant from the US Army Research Office, and by several co-financed ("samfinancierede") PhD scholarships. Recently, center researchers have also started working with several Slovenian researchers and companies on issues in connection with massive terrain data, supported by the Slovenian Research Agency.

As described in Section A2 and A4, the center's multidisciplinary work around terrain data includes co-advising of a Biology PhD student (Moeslund) working on spatial plant diversity modeling. More generally, the center has been a central player in efforts to establish a *Center for Interdisciplinary Geospatial Informatics Research* (CiGIR) at AU, where computer scientists, biologists, geologists and others will collaborate on geospatial modeling questions, such as those arising in biodiversity research. AU (through the University Research Foundation) 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 (Sandel) has already been hired and will spend part of his time at the center. We believe that co-advised PhD students and Post Docs (such as Moeslund and Sandel) are a good way of fulfilling the center's objective of being a catalyst for multidisciplinary collaboration.

Overall, we believe that the center's collaborative efforts and efforts to be a catalyst for multidisciplinary (and industry) research collaboration have been quite successful. Center researchers have collaborated with most, if not all, of the directly relevant algorithms research teams in Denmark (at University of Southern Denmark, the IT-University of Copenhagen, and at the Technical University) and abroad, as well as with selected researchers in other fields. The tradition in the algorithms field for open collaboration and the field's many possibilities for collaboration at conferences, workshops and seminars are obviously an important reason for this. We believe other important, or even essential, drivers of collaboration have been the international strength of the senior research team in the core research areas, as well as the flexibility the center structure provides in terms of funding (which e.g. allows for flexibility in terms of travel, event and visitor funding, and in terms of hiring). As the center has matured and the awareness of the center and its strength has increased, the "MADALGO brand" has probably also been a factor. Finally, the focus on algorithm engineering has been an essential driver for the multidisciplinary (and industry) collaboration. The only major barrier for further (and especially for multidisciplinary and industry) collaboration (as well as center activities in general) has been the limited senior researcher resources at AU. While a majority of the center activities are at AU, only two of the center's six senior researchers are at AU.

B2 Center organization and leadership

Organizationally the center is an integrated part of the Department of Computer Science at AU (with senior faculty Brodal and Arge). Initially, the center also included researchers at MPI (Mehlhorn and Meyer) and MIT (Demaine and Indyk). Around the time of the start of the center Meyer took up a full professor position (chair in algorithm engineering) at Frankfurt University (FRA). In order to keep the research expertise of Meyer and his group in the center, and with the understanding of MPI and the foundation, it was decided to add FRA to the center sites and move the center support for MPI (a 3 year PhD student) to FRA. The center also directly supports activities at MIT (a 9 month PhD student for three years, and a summer month salary for Indyk and Demaine each year). Initially, considerable effort had to be put into negotiations of contracts between the four sites, but other than that the collaboration between the institutions has worked relatively well. As mentioned, the collaboration has led to several concrete research results, has helped to increase the center's visibility and helped to establish a world-leading international environment at AU. However, there have obviously also been challenges in the collaboration, the most significant being the geographical distance between AU and the researchers at MIT, MPI and FRA, which has made it hard to maintain a real sense of belonging to the center outside AU. Although maybe not something that plays a major role, it is also worth noticing that the institutions – especially AU/MADALGO and MPI on the European scene – to some extend are competitors (e.g., when it comes to Post Doc hires). This and the relative proximity (travel time wise) have also naturally resulted in more visits from FRA to AU than from the other sites.

The center is led by a center director (Arge) who is supported by a center manager (Magård) and by a (half time) accountant (Lindstrøm), and with additional secretarial support from the Department of Computer Science. During the center period, it has proven essential to have a full time center manager that is responsible for practical matters such as seminar, workshop and summer school organization, visitor support, maintenance of the center's web-portal, and in general overall coordination of the center's daily activities. Initially, it was proposed that the center director would be assisted by an executive committee consisting of Brodal, Mehlhorn and Indyk. However, during the center period it became clear that an executive committee consisting of Arge and Brodal (as well as Magård) at AU, supported by frequent email discussion and a few meetings each year with all or most of the other senior faculty, was a better solution. With the movement of most of the MPI activities to FRA, Meyer rather than Mehlhorn also became the senior German researcher with most interaction with AU researchers. Finally, the center has made use of an international advisory board, e.g. to provide feedback on the center's annual reports. The board (listed in Appendix a) consists of six world-leading

researchers in the center research areas. Originally, it was planned to also have a Danish advisory board consisting of researchers from Danish research institutions and representatives from Danish industry collaborators in order to support multidiciplinary and industry collaborations. However, we have found that a formal board construction is not needed since ongoing collaborations (as discussed in Section B1 and C1) provide adequate input.

As mentioned, the center is housed in the Department of Computer Science at AU. While it is important for the visibility and "sense of community" of the center that it has its own identity, we also believe that it is important that the center is an integrated part of the Department of Computer Science. Thus in the initial negotiations with the department, the center put considerable emphasis on securing connected facilities with ample space for visitors, meeting, conference and lab facilities, as well as common areas in order to support a dynamic and collaborative environment. Since space is always at a premium, the initial negotiations were somewhat complicated. However, we believe that subsequently the center and the department have mutually benefitted greatly from each other. While the center has been in close proximity to researchers in other related computer science areas and been able to utilize the infrastructure of the department, the center has also impacted the department in a positive way, e.g. by creating (even more) focus on international top-research, by complementing the many applied activities in the department (while at the same time also bridging between theory and practice), and by further improving PhD education in the department (by contributing to an increased PhD student population, PhD classes, and the PhD environment in general).

Overall, we believe the organization and leadership of the center has functioned relatively well. In terms of research and leadership at AU, a generation change had taken place in the algorithms group just prior to the establishment of the center (with two researchers moving to the central faculty administration and Arge and Brodal taking over), and therefore no generation change issues are expected in regard to the future leadership of the center. However as mentioned above, the limited number of senior researchers at AU has to some extend been a limiting factor for center activities. The center has actively been recruiting top senior researchers world-wide, including I/O and cache-oblivious graph algorithms expert Norbert Zeh (for example during his sabbatical stay at the center). In fact, until very recently the plan was for Zeh to join the center in the spring of 2011, but for family reasons he chose to stay in Canada. The center continues its efforts to recruit a senior faculty member that is both highly recognized internationally and complements the research strengths of the existing AU center senior researchers. Efforts are also underway to recruit a highly recognized but more junior faculty member.

C Social Impact of the Center Activities

C1 Industry collaboration

As is relatively normal within computer science, center researchers have collaborated a fair amount with researchers from various industry research laboratories (AT&T, HP, IBM, Google). Through various PhD projects and students, the center has also had informal collaboration with various companies such as BNR (traffic management) and cofman.com (search engines). Center researchers also collaborated with Draper Lab on compressive sensing for astronomical imaging, and play a (minor) role in a technology platform project "A Platform for Galileo based pervasive positioning" that includes several companies (including the Alexandra Institute, Terma and Systematic) along with a range of computer science researchers.

Much of the more intensive center industry collaboration has been on issues in relation to massive terrain data. As mentioned, center researchers have worked with several industry partners, such as COWI and EIVA, as well as with environmental and agricultural researchers, on various terrain problems ranging from removing outliers from raw data to performing water flow and flooding analysis. The work has also been supported by more strategic grants and much of it has been in connection with the TerraStream software package. The center has also had looser collaboration around TerraStream with a number of other companies, including DHI, Statoil, CARIS, and JonesEdmunds. In fact, as a result of the continued purchasing interest from industry collaborators and others, along with the inability of Aarhus University to provide the development and support resources needed to commercialize the software, it was decided during 2009 to establish the company SCALGO (short for Scalable Algorithmics) to commercialize TerraStream. While the company works towards being able to commercialize most of the TerraStream modules, and already has commercialized parts of the outlier detection functionality in collaboration with EIVA, it has also already performed several terrain data processing consulting projects. This includes a project with the Danish Energy Agency (under the Ministry of Climate and

Energy) on assessing the flood risk as a result of rising sea-level using very detailed and massive countrywide terrain data. The result of this project may soon be published on the agency's "klimatilpasning" portal.

One of the center's goals is to be a catalyst for industry collaboration, and in general we think our efforts in that regard have been relatively successful. Our collaboration with industry has not only served to facilitate that relevant research results benefit industry, but it has also influenced the research in the center in a valuable way. It has for example resulted in consideration of several interesting and new research problems, such as the terrain flooding problems discussed in Section A2, just as it has influenced our theoretical model work. As discussed in Section B1, we believe that the focus on algorithm engineering has been an essential driver for our industry collaboration, and that the main barrier for further collaboration has been limited senior faculty resources at AU. Industry collaborations have also benefitted teaching and PhD training in several ways, for example by providing real world examples and data, and by illustrating that algorithm development matters in practice. It is also our experience that algorithms engineering work on real world problems is a good way of introducing new PhD students to the center focus, and in general that a research portfolio consisting of both very theoretical and more practical research problems serves to ensure a PhD student population with diverse interests and strengths.

C2 Dissemination to public

Apart from industry collaboration, the center has tried to disseminate its research and research results to the broader public in Denmark in several ways. Whenever possible the center has conveyed events and results to the general press, which has resulted in more than 40 Danish electronic media, newspaper, magazine and radio/tv article or features, including large features in the science sections of the major Danish newpapers Weekendavisen (twice), Jyllandsposten and Politiken. As mentioned in Section A2, the center's work on the effects of a sea-level rise on Århus also received extensive coverage in the local major newspaper Århus Stiftstidende. Refer to Appendix j. Center researchers have also participated in general Danish public outreach activities. This include activities in connection with robots and internet search engines with primary school children, algorithms activities in connection with visits from or to high school students, as well as presentations on algorithms and MADALGO for high-school teachers (of special math talent classes). They also include presentations at various company events, as well as participation in the yearly Danish research day ("forskningens døgn") with an exhibition on flood prediction. In general, center researchers have given presentations on flood prediction using massive terrain models at several public events.

While it is unclear if the dissemination activities have benefitted the center directly, they have probably done so in several indirect ways. For example, the general public "branding" of the center has undoubtedly created a higher awareness of the center and its activities among Danish research colleagues, students and companies, and therefore probably indirectly helped recruiting efforts and efforts to establish interdisciplinary and industry collaborations – or will do so in the future.

C3 Impact

It is obviously hard to judge which of the center's research findings may have a long lasting impact. Overall, we view the international algorithms research community as the center's most important stakeholder, but the broader research community, industry and society as a whole are also important stakeholders (as also is apparent from the objective of training the next generation of researchers and being a catalyst for collaboration). In terms of the algorithms community, a number of the center's concrete research results have the potential for having a lasting impact (e.g., our I/O, cache-oblivious and internal memory model results on higher-dimensional range reporting, our results on general streaming techniques, and our massive terrain and sorting engineering results), but a more high-level long lasting impact might be the reinforcement of the community's focus on developing and investigating more realistic algorithmic models for massive data processing, as well as on the interaction between theoretical developments and algorithm engineering. Hopefully the center's multidisciplinary work with especially environmental researchers will also have a lasting impact. In terms of the impact on the Danish society at large, the center's long lasting impact is likely to be found in the training and educational activities, especially in the training of computer scientists in an interdisciplinary environment with focus on both theoretical and practical issues. Through the center's spin-off SCALGO and collaboration with industry and government authorities, the center's work on software for processing massive terrain data, and for water flow modeling and flood prediction, is also likely to have a lasting impact.

Appendices

Appendix a through i contain tables with information as specified in the foundation guidelines. Appendix j contains information about public outreach in the foundation format used in the annual center reports, and Appendix k contains a list of 10 selected center publications.
Note that much of the information in the appendices is *only* given for AU employees (in which case this is specified in notes below the tables). Note also that due to its volume, some of the information may be incomplete.

List of appendices:

a) Senior staff associated with center
b) Post docs affiliated with center
c) PhD students supervised within center
d) Foreign scientists (and PhD students) spending more than 3 weeks at center
e) International conferences, symposia and workshops organized by the center
f) Publications
g) Patents and applications
h) Other relevant information documenting international standing of center
    - Awards and acknowledgments
    - AU external funding
    - AU editorial work
    - Center conference and seminar contributions
    - Center collaboration
i) Information on center staffs involvement in educational activities
j) Public outreach
k) Selected center publications

a) Senior staff associated with center

AU

Name	Position	Nationality	Period	Funded by (share of full time equivalent)[1]		
				DNRF	AU	Other
Lars Arge	Professor	Denmark	All period	1/10	6/10	3/10
Gerth S. Brodal	Associate Professor	Denmark	All period		1	

MIT

Name	Position	Nationality	Period	Funded by (share of full time equivalent)		
				DNRF	AU	Other
Piotr Indyk	Associate Professor	Poland	2 months a year	1/2		1/2
Erik Demaine	Associate Professor	USA	2 months a year	1/2		1/2

Note: The center funds one month summer salary for Demaine and Indyk in each of the 5 center years.

MPI/FRA

Name	Position	Nationality	Period	Funded by (share of full time equivalent)		
				DNRF	AU	Other
Ulrich Meyer	Professor	Germany	3 months a year			1
Kurt Mehlhorn	Professor	Germany	1 month a year			1

Note:

- During the first center year, Meyer moved from MPI to Frankfurt University.
- The center is advised by an international advisory board consisting of world-leading researchers in the center research areas:
    - Pankaj K. Agarwal, Chair and RJR Nabisco Professor of Computer Science at Duke University
    - Giuseppe F. Italiano, Professor of Computer Science and Chair of the Department of Computer Science, Systems and Production at University of Rome "Tor Vergata"
    - J. Ian Munro, University Professor and Canada Research Chair in Algorithms Design in the School of Computer Science at the University of Waterloo
    - S. Muthu Muthukrishnan, Research Scientist at Google New York
    - Peter Sanders, Professor of Computer Science at Karlsruhe University
    - Jeff S. Vitter, Provost and executive vice chancellor and Professor of Computer Science at the University of Kansas

[1] In all tables, funding information indicates relative (and approximate) distribution between funding sources (Danish National Research Foundation, Aarhus University and Other) regardless of hiring period.

b) Post docs affiliated with center

AU

Name	Period	Nationality	PhD from	Present employment	Funded by (share of full time equivalent)		
					DNRF	AU	Other
Henrik Blunk	-10.08	Germany	Munster 2006	Post doc Aarhus			1
Srinivas Rao	08.07- 01.09	India	Chennai 2002	Assist. Prof. Seoul National	1		
Mohammad Abam	01.08-01.10	Iran	Eindhoven 2007	Post Doc Dortmund	1		
Peter Hachenberger	07.08-06.09	Germany	MPI 2006	German Industry	1		
Deepak Ajwani	10.08-09.10	India	MPI 2008		1		
Peyman Afshani	02.09-08.10	Iran	Waterloo 2008		2/3		1/3
Nodari Sitchinava	08.09-07.11	Georgia	UC Irvine 2009		1		
Brody Sandel	06.10-(05.11)	USA	Berkeley 2010			1	
Elad Verbin	08.10-(07.11)	Israel	Tel-Aviv 2007		1/2		1/2
Qin Zhang	08.10-(07.11)	China	Hong Kong 2010		1		

Notes:

- Ajwani and Afshani will leave the center this fall. Ajwani is likely to join Center for Unified Computing at Cork University on a grant from IBM and the Irish Research Council. Afshani has reviewed one of the prestigious 2 year postdoctoral fellowships awards from the Natural Sciences and Engineering Research Council of Canada (NSERC). He will spend the fall at Dalhousie University (with Norbert Zeh) and may then return to the center.
- The three last (and newly hired) Post Docs are initially hired for one year but with possibility for extension (thus end dates are in parenthesis). The funding distribution is based on the first year funding. To insure continued development in research focus and goals, the center normally restricts Post Doc appointments to two years.

MPI/FRA

Name	Period	Nationality	PhD from	Present employment	Funded by (share of full time equivalent)		
					DNRF	AU	Other
Sathish Govindarajan	-09.07	India	Duke 2004	Assist. Prof. IICS India			1
Kevin Chang	-09.08	USA	Yale 2007	US Industry			1
Gabriel Moruz	10.07-(03.11)	Romania	Aarhus 2007				1

c) PhD students supervised within center

AU

Name	Period	Nationality	Year of degree	Present employment	Funded by (share of full time equivalent)		
					DNRF	AU	Other
Johan Nilsson	-07.07	Sweden	2007	Swedish Industry		1	
Gabriel Moruz	-07.07	Romania	2007	Post Doc Frankfurt		1	
Anders H. Jensen	-07.08	Denmark	-	Danish Industry			1
Martin Olsen	-07.09	Denmark	2009	Assoc. Prof. AU, Herning		1/3	2/3
Thomas Mølhave	-08.09	Denmark	2009	Post Doc Duke		1/3	2/3
Allan G. Jørgensen	-01.10	Denmark	2010	Danish Industry	1/3		2/3
Lasse Deleuran	08.07-(07.11)	Denmark	(2011)		1/3	1/3	1/3
Kostas Tsakalidis	08.07-(07.11)	Greece	(2011)		2/3	1/3	
Jesper E. Moeslund	02.08-(01.12)	Denmark	(2012)		1/6	1/3	3/6
Mark Greve	02.08-(01.12)	Denmark	(2012)		2/3	1/3	
Morten Revsbæk	02.08-(01.13)	Denmark	(2013)		1/5	2/5	2/5
Pooya Davoodi	05.08-(04.11)	Iran	(2011)		1		
Jacob Truelsen	08.08-(07.13)	Denmark	(2013)		1/2	1/2	
Kasper D. Larsen	08.08-(07.13)	Denmark	(2013)		1/6	5/6	
Casper Kejlberg-Rasmussen	02.09-(07.13)	Denmark	(2013)		1/3	2/3	
Freek van Walderveen	08.09-(07.12)	Netherlands	(2012)		1/6	2/3	1/6
Sarah Zakarias	08.09-(07.13)	Denmark	(2013)				1

Note:

- PhD students at AU are normally admitted in the so-called *4+4 PhD program* after four years of university study. The PhD program itself then consists of a two year part A and a two year part B, with a qualification exam (where a master degree is also obtained) between the two parts. In order to recruit (and thus retain) highly talented students early in their study, and to potentially better recruitment possibilities for skilled bachelor students from other institutions, the Faculty of Science recently introduced a "honors" program, or the so-called *3+5 program*, where PhD students are admitted already after 3 years of university study and then enter the 4+4 program after a year where different PhD subjects can be explored. Finally, students with a master degree can be admitted in the three year so-called *5+3 PhD program*.
- Center students Olsen, Revsbæk, Davoodi and van Walderveen were admitted in the 5+3 program and students Larsen, Kejlberg-Rasmussen and Zakarias in the 3+5 program. The rest of the AU center students were admitted in the 4+4 program. Dates in parenthesis indicated estimated finish times based on admission time and program.
- Center student Jensen left the 4+4 PhD program after part A with a master degree, and 3+5 honor student Zakarias is likely to choose an advisor outside the center as she enters the 4+4 program.
- Center student Revsbæk will soon return after a one year leave of absence to work in the center spin-off company SCALGO. After finishing his qualification exam at the end of July 2010, center student Truelsen has taken a one year leave of absence to work at SCALGO.
- Arge was/is advisor for Jensen, Mølhave, Deleuran, Moeslund (co-advisor), Revsbæk, Larsen, van Walderveen and Zakarias. Brodal was/is for Nilsson, Moruz, Olsen, Jørgensen, Tsakalidis, Greve, Davoodi, Truelsen and Kejlberg-Rasmussen.

- The funding information indicates relative distribution of the total (*also future*) expenses for each PhD student, where the cost of the first year for 3+5 students (paid by AU) is ignored and where the cost of part A and part B of 4+4 students are estimated at 1/3 and 2/3 of the total cost, respectively.

MPI/FRA

Name	Period	Nationality	Year of degree	Present employment	Funded by (share of full time equivalent)		
					DNRF	AU	Other
Deepak Ajwani	-09.08	India	2008	Post Doc Aarhus Uni.			1
Andreas Beckmann	All period	Germany	(2011)		5/6		1/6
Andrei Negoescu	01.09-	Romania	(2012)				1
Volker Weichert	01.10-	Germany	(2013)				1

Note:
- Meyer was/is advisor or co-advisor for all students.
- According to the contract with Frankfurt University, the center will fund one PhD student for three years at Frankfurt University. All other center expenses at Frankfurt are covered by other sources. The funding information indicates relative distribution of PhD student expenses until now.

MIT

Name	Period	Nationality	Year of degree	Present employment	Funded by (share of full time equivalent)		
					DNRF	AU	Other
Mihai Patrascu	-05.08	Romania	2008	AT&T Research	5/8		3/8
Anastasio Sidiropoulos	-05.08	Greece	2008	Post doc Toronto			1
Oren Weimann	-03.09	Israel	2009	Post doc Weizmann			1
Khan Do Ba	All period	Vietnam	(2012)		1/3		2/3
Jelani Nelson	All period	USA	(2011)				1
Eric Price	08.09-	USA	(2013)				1
Morteza Zadimoghaddam	08.09-	Iran	(2013)		1/2		1/2

Note:
- Demaine was/is advisor for Patrascu, Weimann and Zadimoghaddam. Indyk was/is advisor for Sidiropoulos, Do Ba and Price. Nelson is co-advised by Demaine and Indyk.
- According to the contract with MIT, the center will fund nine months of one PhD student at MIT in each of three years. Apart from summer salary for Demaine and Indyk, all other center expenses at MIT are covered by other sources. The funding information indicates relative distribution of PhD student expenses until now.

d) Foreign scientists (and PhD students) spending more than 3 weeks at center

Name	Position	Period	Nationality	Employment	Funded by (share of full time equivalent)		
					DNRF	AU	Other
Bradford Nickerson	Professor	03.07-06.08	Canada	New Brunswick			1
Michael T. Goodrich	Professor	08.07-08.07	USA	UC Irvine			1
Nodari Sitchinava	PhD student	08.07-08.07	Georgia	UC Irvine			1
Igor Nitto	PhD student	08.07-09.07	Italy	Pisa			1
Jeremy Fineman	PhD student	06.08-06.08	USA	MIT			1
Muriel Dulieu	PhD student	07.08-08.08	Belgium	Polytechnic Inst. of NYU			1
John Iacono	Associate Professor	07.08-08.08	USA	Polytechnic Inst. of NYU	1/2		1/2
Nodari Sitchinava	PhD student	08.08-09.08	Georgia	UC Irvine	1/2		1/2
Andrew Danner	Assistant Professor	08.08-08.08	USA	Swarthmore College	1		
Jan van Leeuwen	Professor	05.09-06.09	Netherlands	Utrecht			1
Martin Šmérek	PhD Student	08.09-11.09	Czech	Masaryk			1
Norbert Zeh	Associate Professor	08.09-02.10	Germany	Dalhousie			1
Shervin Daneshpajouh	PhD Student	01.10-08.10	Iran	Sharif			1
John Iacono	Associate Professor	08.10-08.10	USA	Polytechnic Inst. of NYU	1		

Note:

- Table *only* includes visits to AU by non-center researchers and students.
- Funding information only relates to salary and honorariums. In general, visitors tend to prefer getting housing and travel expenses reimbursed, which the center did for several of the listed visitors.
- Center has also hosted many shorter term visitors (several for more than one period of time), including faculty Athanasios Tsakalidis (Patras), Peter Sanders (Karlsruhe), S. Muthu Muthukrishnan (Google and Rutgers), Charles Leiserson (MIT), Jeff Vitter (Kansas), Pino Italiano (Rome), Martin Strauss (Michigan), Sudipto Guha (Penn), Ravi Kumar (Yahoo!), T.S. Jayram (IBM Almaden), Jonathan Shewchuk (Berkeley), Philip Bille (ITU), Rolf Fagerberg (Southern Denmark), Rajeev Raman (Leicester), Herman Haverkort (Eindhoven), Alex Lopez-Ortiz (Waterloo), Ian Munro (Waterloo), Seth Pettie (Michigan), Alexander Wolff (Eindhoven), Mikkel Thorup (AT&T), Spyros Sioutas (Ionian), Kostas Tsichlas (Thessaloniki), Jérémy Barbay (Chile) and Rasmus Pagh (ITU), Sariel Har-Peled (UIUC), Jack Snoeyink (UNC), Martin Isenburg (LLNL), Bardia Sadri (Toronto), Ke Yi (HKUST), Riko Jacob (Munich), Christian Knauer (FU Berlin), Morteza Monemizadeh (Dortmund), Mihai Patrascu (AT&T), Timothy Chan (Waterloo), Norbert Zeh (Dalhausie), and Hamish Carr (Leeds), as well as students Arash Farzan (Waterloo), Dmitriy Morozov (Duke), Jeff Philips (Duke), Ana Krulec (Primorska), Jionqxin Jin (HKUST) and Man Kwun Chiu (HKUST).,

e) International conferences, symposia and workshops organized by the center

Title and date of event	Number of participants	
	Danish	International
MADALGO Inauguration, August 24, 2007	~50	~40
MADALGO Summer school on data stream algorithms, August 20-23, 2007	~25	~45
Dagstuhl Seminar on Structure Theory and FPT Algorithmics for Graphs, Digraphs and Hypergraphs, July 8-13, 2007	2	~65
MADALGO Summer school on cache-oblivious algorithms, August 18-21, 2008	30	30
Dagstuhl seminar on Data Structures, February 17-22, 2008	4	45
Workshop on Massive Data Algorithmics (MASSIVE), June 11, 2009	25	38
Symposium on Computational Geometry (SoCG), June 8-10, 2009	26	110
25th Symposium on Computational Geometry Celebration, June 7, 2009	20	98
Workshop on Massive Data Algorithmics (MASSIVE), June 17, 2010	6	24
Dagstuhl seminar on Data Structures, February 28 – March 3, 2010	5	45
MADALGO Summer School on Geometric Data Structures, August 16-19, 2010	23	28

Note:

- The listed Dagstuhl seminars are seminars where one of the senior center staff was part of the organizing committee.
- In addition to the above events, the center arranged a large number of smaller events, including a number of informal workshops around the TerraStream software.

f) Publications

Conference proceedings

	Year	Authors	Title	Venue	
C1	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)
C2	2007	S. Guha, P. Indyk and A. McGregor	Sketching Information Divergences	Proc. Annual Conference on Learning Theory (COLT)	(PR)(CO)
C3	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)
C4	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)
C5	2007	A. G. Jørgensen, G. Moruz and T. Mølhave	Resilient Priority Queues	Proc. International Workshop on Algorithms and Data Structures (WADS)	(PR)
C6	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)
C7	2007	A. Danner, T. Mølhave, K. Yi, P. K. Agarwal, L. Arge and H. Mitasova	TerraStream: From Elevation Data to Watershed Hierarchies	Proc. ACM International Symposium on Advances in Geographical Information Systems (ACM-GIS)	(PR)(CO)
C8	2007	M. Patrascu and Mikkel Thorup	Planning for Fast Connectivity Updates	Proc. IEEE Symposium on Foundations of Computer Science (FOCS)	(PR)(CO)
C9	2007	G. Franceschini, S. Muthukrishnan, and M. Patrascu	Radix Sorting With No Extra Space	Proc. European Symposium on Algorithms (ESA)	(PR)(CO)
C10	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)
C11	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)
C12	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)
C13	2007	M. Patrascu	Lower Bounds for 2-Dimensional Range Counting	Proc. ACM Symposium on Theory of Computing (STOC)	(PR)
C14	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)
C15	2007	T. Friedrich and D. Ajwani	Average-Case Analysis of Online Topological Ordering	Proc. International Symposium on Algorithms and Computation (ISAAC)	(PR)

C16	2007	K. Chang	Multiple pass streaming algorithms for learning mixtures of distributions in R^d	Proc. Algorithmic Learning Theory (ALT)	(PR)
C17	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)
C18	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)
C19	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)
C20	2007	M. Olsen	Nash Stability in Additively Separable Hedonic Games is NP-hard	Proc. Conference on Computability in Europe (CiE)	(PR)
C21	2008	M. Ruzic and P. Indyk	Near-Optimal Sparse Recovery in the L1 norm	Proc. Symposium on Foundations of Computer Science (FOCS)	(PR)(CO)
C22	2008	M. Patrascu	(Data) STRUCTURES	Proc. Symposium on Foundations of Computer Science (FOCS)	(PR)
C23	2008	M. Patrascu	Succincter	Proc. Symposium on Foundations of Computer Science (FOCS)	(PR)
C24	2008	E. Demaine, S. Langerman and E. Price	Confluently Persistent Tries for Efficient Version Control	Proc. Scandinavian Workshop on Algorithm Theory (SWAT)	(PR)(CO)
C25	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)
C26	2008	U. Meyer	On Dynamic Breadth-First Search in External-Memory	Proc. Symposium on Theoretical Aspects (STACS)	(PR)
C27	2008	U. Meyer	On Trade-Offs in External-Memory Diameter Approximation	Proc. Scandinavian Workshop on Algorithm Theory (SWAT)	(PR)
C28	2008	G. S. Brodal and A. G. Jørgensen	Selecting Sums in Arrays	Proc.International Symposium on Algorithms and Computation (ISAAC)	(PR)
C29	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)
C30	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)
C31	2008	M. Olsen	The Computational Complexity of Link Building	Proc. International Conference on Computing and Combinatorics (COCOON)	(PR)
C32	2008	M.A. Abam, M. de Berg and J. Gudmundsson	A Simple and Efficient Kinetic Spanner	Proc. Symposium on Computational Geometry (SoCG)	(PR)(CO)
C33	2008	L. Arge, M.T. Goodrich, M. Nelson and N. Sitchinava	Fundamental Parallel Algorithms for Private-Cache Multiprocessors	Proc. Symposium on Parallelism in Algorithms and Architectures (SPAA)	(PR)(CO)

C34	2008	L. Arge, T. Moelhave and N. Zeh	Cache-Oblivious Red-Blue Line Segment Intersection	Proc. European Symposium on Algorithm (ESA)	(PR)(CO)
C35	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)
C36	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)
C37	2008	P. Indyk	Explicit Constructions for Compressed Sensing of Sparse Signals	Proc Symposium on Discrete Algorithms (SODA)	(PR)
C38	2008	A. Andoni, P. Indyk and R. Krauthgamer	Earth Mover Distance over High-Dimensional Spaces	Proc. Symposium on Discrete Algorithms (SODA)	(PR)(CO)
C39	2008	P. Indyk and A. McGregor	Declaring Independence via the Sketching of Sketches	Proc. Symposium on Discrete Algorithms (SODA)	(PR)(CO)
C40	2008	K. Onak and A. Sidiropoulos	Circular Partitions with Applications to Visualization and Embeddings	Proc. Symposium on Computational Geometry (SoCG)	(PR)(CO)
C41	2008	J. Matousek and A. Sidiropoulos	Inapproximability for metric embeddings into R^d	Proc. Symposium on Foundations of Computer Science (FOCS)	(PR)(CO)
C42	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)
C43	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)
C44	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)
C45	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)
C46	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)
C47	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)
C48	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)
C49	2008	E. Demaine, M. Badoiu, 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)
C50	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)
C51	2008	E. R. Hansen, S. S. Rao and P. Tiedemann	Compressing Binary Decision Diagrams	European Conference on Artificial Intelligence (ECAI)	(PR)(CO)

C52	2008	R. Berinde, P. Indyk and M. Ruzic	Practical Near-Optimal Sparse Recovery in the L1 Norm (invited paper)	Proc. Allerton Conference	(CO)
C53	2008	R. Berinde, A. Gilbert, P. Indyk, H. Karloff and M. Strauss	Combining Geometry and Combinatorics: A Unified Approach to Sparse Signal Recovery	Proc. Allerton Conference	(CO)
C54	2008	M.A. Abam, M. de Berg, and S-H. Poon	Fault-Tolerant Conflict-Free Coloring	Proc. Canadian Conference on Computational Geometry	(CO)
C55	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)
C56	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)
C57	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)
C58	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)
C59	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)
C60	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)
C61	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)
C62	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)
C63	2009	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	
C64	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)
C65	2009	A. Beckmann, R. Dementiev and J. Singler	Building a Parallel Pipelined External Memory Algorithm Library	Proc. International Symposium on Parallel and Distributed Processing (IPDPS)	(PR)
C66	2009	G. S. Brodal and A. Jørgensen	Data Structures for Range Median Queries	Proc. International Symposium on Algorithms and Computation (ISAAC)	(PR)
C67	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)
C68	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)
C69	2009	G. S. Brodal, A. Jørgensen and T. Mølhave	Fault Tolerant External Memory Algorithms	Proc. Algorithms and Data Structures Symposium (WADS)	(PR)

C70	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)
C71	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)
C72	2009	M. Abam and M. de Berg	Kinetic Spanners in R^d	Proc. Symposium on Computational Geometry (SoCG)	(PR)(CO)
C73	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)
C74	2009	D. Ajwani	On P-complete Problems in Memory Hierarchy Models	Proc. Workshop on Massive Data Algorithmics (MASSIVE)	
C75	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)
C76	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)
C77	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)
C78	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)
C79	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)
C80	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)
C81	2009	P. Afshani, C. Hamilton and N. Zeh	A General Approach for Cache-Oblivious Range Reporting and Approximate Range Counting	Proc. Symposium on Computational Geometry (SoCG)	(PR)(CO)
C82	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)
C83	2009	P. Afshani, J. Barbay and T. Chan	Instance-optimal Geometric Algorithms	Proc Symposium on Foundations of Computer Science (FOCS)	(PR)(CO)
C84	2009	L. Arge, M.T. Goodrich and N. Sitchinava	Parallel External Memory Model	Proc. Workshop on Theory and Many-Cores	
C85	2009	L. Arge and M. Revsbæk	I/O-Efficient Contour Tree Simplification	Proc. International Symposium on Algorithms and Computation (ISAAC)	(PR)
C86	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)

C87	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)
C88	2009	R. Berinde and P. Indyk	Sequential Sparse Matching Pursuit	Proc. Allerton Conference	(PR)(CO)
C89	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)
C90	2009	A. Andoni, P. Indyk, K. Onak and R. Rubinfeld	External Sampling	Proc. International Colloquium on Automata, Languages and Programming (ICALP)	(PR)(CO)
C91	2009	E. Demaine, M. Demaine, G. Konjevod and R. Lang	Folding a Better Checkerboard	Proc. International Symposium on Algorithms and Computation (ISAAC)	(PR)(CO)
C92	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)
C93	2009	E. Demaine, M. Hajiaghayi, and D. Marx	Minimizing Movement: Fixed-Parameter Tractability	Proc. European Symposium on Algorithms (ESA)	(PR)(CO)
C94	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)
C95	2009	E. Demaine, D. Kane and G. Price	A Pseudopolynomial algorithm for Alexandrov's Theorem	Proc. Algorithms and Data Structures Symposium (WADS)	(PR)(CO)
C96	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)
C97	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)
C98	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)
C99	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)
C100	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)
C101	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)
C102	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)
C103	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)

C104	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)
C105	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)
C106	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)
C107	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)
C108	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)
C109	2009	P. Klein, S. Mozes and O. Weimann	Shortest Paths in Directed Planar Graphs with Negative Lengths: A Linear-Space O(n log n)-Time Algorithm	Proc. Symposium on Discrete Algorithms (SODA)	(PR)(CO)
C110	2010	K. Do Ba, P. Indyk, E. Price and D.P. Woodruff	Lower Bounds for Sparse Recovery	Proc. Symposium on Discrete Algorithms (SODA)	(PR)(CO)
C111	2010	P. Indyk, H.Q. Ngo and A. Rudra	Efficiently Decodable Non-adaptive Group Testing	Proc. Symposium on Discrete Algorithms (SODA)	(PR)(CO)
C112	2010	D.M. Kane, J. Nelson and D.P. Woodruff	An Optimal Algorithm for the Distinct Elements Problem	Proc. Symposium on Principles of Database Systems (PODS)	(PR)(CO)
C113	2010	J. Nelson and D.P. Woodruff	Fast Manhattan Sketches in Data Streams	Proc. Symposium on Principles of Database Systems (PODS)	(PR)(CO)
C114	2010	I. Diakonikolas, D.M. Kane and J. Nelson	Bounded Independence Fools Degree-2 Threshold Functions	Proc. Symposium on Foundations of Computer Science (FOCS)	(PR)(CO)
C115	2010	D.M. Kane, J. Nelson and D.P. Woodruff	On the Exact Space Complexity of Sketching and Streaming Small Norms	Proc. Symposium on Discrete Algorithms (SODA)	(PR)(CO)
C116	2010	A. Beckmann, U. Meyer, P. Sanders and J. Singler	Energy-Efficient Sorting using Solid State Disks	Proc. International IEEE Green Computing Conference	(PR)(CO)
C117	2010	M. Greve, A.G. Jørgensen, K.D. Larsen and J. Truelsen	Cell Probe Lower Bounds and Approximations for Range Mode	Proc. International Colloquium on Automata, Languages and Programming (ICALP)	(PR)
C118	2010	M. Olsen	Maximizing PageRank with new Backlinks	Proc. International Conference on Algorithms and Complexity (CIAC)	(PR)
C119	2010	G.S. Brodal, E. Demaine, J. T. Fineman, J. Iacono, S. Langerman and J.I. Munro	Cache-Oblivious Dynamic Dictionaries with Optimal Update/Query Tradeoff	Proc. Symposium on Discrete Algorithms (SODA)	(PR)(CO)
C120	2010	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)
C121	2010	M.A. Abam and S. Har-Peled	New constructions of SSPDs and their applications	Proc. Symposium on Computational Geometry (SoCG)	(PR)(CO)

C122	2010	M.B. Kjærgaard, H. Blunck, T. Godsk, T. Toftkjær, D.L. Christensen, and K. Grønbæk	Indoor Positioning using GPS Revisited	Proc. International Conference on Pervasive Computing (Pervasive)	(PR)
C123	2010	L. Arge, M.T. Goodrich and N. Sitchinava	Parallel external memory graph algorithms	Proc. International Parallel & Distributed Processing Symposium (IPDPS)	(PR)(CO)
C124	2010	P. Afshani, L. Arge and K.D. Larsen	Orthogonal Range Reporting: Query Lower Bounds, Optimal Structures in 3-d, and Higher Dimensional Improvements	Proc. Symposium on Computational Geometry (SoCG)	(PR)
C125	2010	P. Afshani, L. Arge and K.D Larsen	I/O-Efficient Orthogonal Range Reporting in Three and Higher Dimensions	Proc. Workshop on Massive Data Algorithmics (MASSIVE)	
C126	2010	T. Mølhave, P.K. Agarwal, L. Arge and M. Revsbæk	Scalable Algorithms for Large High-Resolution Terrain Data	Proc. International Conference on Computing for Geospatial Research & Application (COM.GEO)	(PR)(CO)
C127	2010	L. Arge, M. Revsbæk and Norbert Zeh	I/O-Efficient Computation of Water Flow Across a Terrain	Proc. Symposium on Computational Geometry (SoCG)	(PR)(CO)
C128	2010	G.S. Brodal, P. Davoodi and S.S. Rao	On Space Efficient Two Dimensional Range Minimum Data Structures	Proc. European Symposium on Algorithms (ESA)	(PR)(CO)
C129	2010	D. Ajwani, N. Sitchinava and N. Zeh	Geometric Algorithms for Private-Cache Chip Multiprocessors	Proc. European Symposium on Algorithms (ESA)	(PR)(CO)
C130	2010	Z. Abel, N. Benbernou, M. Damian, E.D. Demaine, M.L. Demaine, R. Flatland, S. Kominers and R. Schwelle	Shape Replication Through Self-Assembly and RNase Enzymes	Proc. Symposium on Discrete Algorithms (SODA)	(PR)(CO)
C131	2010	E.D. Demaine, M. Hajiaghayi and K. Kawarabayashi	Decomposition, Approximation, and Coloring of Odd-Minor-Free Graphs	Proc. Symposium on Discrete Algorithms (SODA)	(PR)(CO)
C132	2010	N. Gershenfeld, D. Dalrymple, K. Chen, A. Knaian, F. Green, E.D. Demaine, S. Greenwald and P. Schmidt-Nielsen	Reconfigurable Asynchronous Logic Automata	Proc. Symposium on Principles of Programming Langauges (POPL)	(PR)(CO)
C133	2010	G. Aloupis, J. Cardinal, S. Collette, E.D. Demaine, M.L. Demaine, M. Dulieu, R. Fabila-Monroy, V. Hart, F. Hurtado, S. Langerman, M. Saumell, C. Seara and P. Taslakian	Matching Points with Things	Proc. Latin American Theoretical Informatics Symposium (LATIN)	(PR)(CO)
C134	2010	E.D. Demaine and M. Zadimoghaddam	Scheduling to Minimize Power Consumption using Submodular Functions	Proc. Symposium on Parallelism in Algorithms and Architectures (SPAA)	(PR)
C135	2010	S. Gilbert, R. Guerraoui, F. Malakouti and M. Zadimoghaddam	Collaborative Scoring with Dishonest Participants	Proc. Symposium on Parallelism in Algorithms and Architectures (SPAA)	(PR)(CO)
C136	2010	N. Alon, E.D. Demaine, M. Hajiaghayi and T. Leighton	Basic Network Creation Games	Proc. Symposium on Parallelism in Algorithms and Architectures (SPAA)	(PR)(CO)
C137	2010	E.D. Demaine and M. Zadimoghaddam	Minimizing the Diameter of a Network using Shortcut Edge	Proc. Scandinavian Workshop on Algorithm Theory (SWAT)	(PR)

C138	2010	M. Bateni, M.H. Hajiaghayi and M. Zadimoghaddam	Submodular Secretary Problem and Extensions	Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX)	(PR)(CO)

Journals

	Year	Authors	Title	Venue	
J1	2007	G. S. Brodal, R. Fagerberg and G. Moruz	On the Adaptiveness of Quicksort	ACM Journal of Experimental Algorithmics, 12	(PR) (CO)
J2	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)
J3	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)
J4	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)
J5	2009	M. Olsen	Nash Stability in Additively Separable Hedonic Games and Community Structures	Theory of Computing Systems, 45(4)	(PR)
J6	2009	M. Abam, M. de Berg, M. Farshi and J. Gudmundsson	Region-Fault Tolerant Geometric Spanners	Discrete & Computational Geometry, 41(4)	(PR)(CO)
J7	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)
J8	2009	L. Arge, V. Samoladas and K. Yi	Optimal External-Memory Planar Point Enclosure	Algorithmica, 54(3)	(PR)(CO)
J9	2009	L. Arge, M. de Berg and H. Haverkort	Cache-Oblivious R-Trees	Algorithmica, 53(1)	(PR)(CO)
J10	2009	H. Iben, J. O'Brien and E. Demaine	Refolding Planar Polygons	Discrete & Computational Geometry, 41(3)	(PR)(CO)
J11	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)
J12	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)
J13	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)
J14	2010	P. Indyk and A. Gilbert	Sparse Recovery Using Sparse Matrices	Proceedings of the IEEE June 2010	(PR)(CO)
J15	2010	E.D. Demaine, S.Langerman and E. Price	Confluently Persistent Tries for Efficient Version Control	Algorithmica 57(3)	(PR)(CO)
J16	2010	M.A. Abam, M. de Berg, P. Hachenberger and A. Zarei	Streaming Algorithms for Line Simplification	Discrete & Computational Geometry 43(3)	(PR)(CO)
J17	2010	M.A. Abam, M. de Berg and J. Gudmundsson	A Simple and Efficient Kinetic Spanner	Computational Geometry: Theory and Applications 43(3)	(PR)(CO)

J18	2010	D. Ajwani and T. Friedrich	Average-case Analysis of Incremental Topological Ordering	Discrete Applied Mathematics 158	(PR)(CO)
J19	2010	H. Blunck and J. Vahrenhold	In-Place Algorithms for Computing (Layers of) Maxima	Algorithmica 57(1)	(PR)(CO)
J20	2010	P. Indyk, Z. Syed, C. Stultz, M. Kellis and J. Guttag	Motif discovery in physiological datasets: A methodology for inferring predictive elements	ACM Transactions on Knowledge Discovery in Data 4(1)	(PR)(CO)
J21	2010	E. Hawkes, B. An, N. M. Benbernou, H. Tanaka, S. Kim, E.D. Demaine, D. Rus and R.J. Wood	Programmable matter by folding	Proceedings of the National Academy of Sciences of the United States of America 107(28)	(PR)(CO)
J22	2010	J.L. Bredin, E.D. Demaine, M. Hajiaghayi and D. Rus	Deploying Sensor Networks with Guaranteed Fault Tolerance	IEEE/ACM Transactions on Networking 18(1)	(PR)(CO)
J23	2010	E.D. Demaine, J. Iacono and S. Langerman	Grid Vertex-Unfolding Orthostacks	International Journal of Computational Geometry and Applications 20(3)	(PR)(CO)

Thesis

	Year	**Authors**	**Title**	**Cite**	
T1	2007	I. Brudaru	Heuristics for Average Diameter Approximation with External Memory Algorithms	MPI	MS Thesis
T2	2007	G. Moruz	Hardware-Aware Algorithms and Data Structures	AU	PhD Thesis
T3	2008	M. Patrascu	Lower Bound Techniques for Data Structures	MIT	PhD Thesis
T4	2008	A. Sidiropoulos	Computational metric embeddings	MIT	PhD Thesis
T5	2008	D. Ajwani	Traversing large graphs in realistic settings	MPI	PhD Thesis
T6	2008	K. Do Ba	Testing closeness of distributions under the EMD metric	MIT	MS Thesis
T7	2008	K. Lai	Complexity of Union-Split-Find Problems	MIT	MS Thesis
T8	2008	J. M. Larsen and M. Nielsen	En undersøgelse af algoritmer til løsning af generalized movers problem i 3D	AU	MS Thesis
T9	2008	C. Andersen	An optimal minimum spanning tree algorithm	AU	MS Thesis
T10	2008	M. Revsbæk	I/O-efficient Algorithms for Batched Union-Find with Dynamic Set Properties and its Applications to Hydrological Conditioning	AU	MS Thesis
T11	2008	A. H. Jensen	I/O-efficient Processing of LIDAR Data	AU	MS Thesis
T12	2009	Martin Olsen	Link Building	AU	PhD Thesis
T13	2009	Thomas Mølhave	Handling Massive Terrains and Unreliable Memory	AU	PhD Thesis

T14	2009	Henrik B. Kirk	Searching with Dynamic Optimality: In Theory and Practice	AU	MS Thesis
T15	2009	Krzysztof Piatkowski	Implementering og udvikling af maksimum delsum algoritmer	AU	MS Thesis
T16	2009	O. Weimann	Accelerating Dynamic Programming	MIT	PhD Thesis
T17	2009	Volker Weichert	Radiation parameterization of the climate model COSMO/CLM in CUDA	FRA	MS Thesis
T18	2009	R. Berinde	Advances in Sparse Signal Recovery Methods	MIT	MS Thesis
T19	2010	A. G. Jørgensen	Data Structures: Sequence Problems, Range Queries, and Fault Tolerance	AU	PhD Thesis
T20	2010	J. Moeslund	Fine-resolution geospatial modeling of contemporary and potential future plant diversity in Denmark	AU	MS Thesis
T21	2010	J. Truelsen	Working Set Implicit Dictionaries and Range Mode Lower Bounds and Approximations	AU	MS Thesis

Other

	Year	**Authors**	**Title**	**Venue**	
O1	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)
O2	2008	A. Andoni and P. Indyk	Near-Optimal Hashing Algorithms for Approximate Nearest Neighbor in High Dimensions	Communications of the ACM, 51(1)	(CO)
O3	2008	K. Mehlhorn and P. Sanders	Algorithms and Data Structures: The Basic Toolbox	Springer Verlag	(CO)
O4	2009	D. Ajwani and U. Meyer	Design and Engineering of External Memory Traversal Algorithms for general graphs	In Algorithmic of Large and Complex Networks, Springer Verlag	(PR)
O5	2009	L. Arge and N. Zeh	External-memory Algorithms and Data Structures	In Algorithms and Theory of Computation Handbook, CRC Press	(PR)(CO)
O6	2009	R. Hearn and E. Demaine	Games, Puzzles, and Computation	A.K. Peters	(CO)
O7	2010	D. Ajwani and H. Meyerhenke	Realistic Computer Models	In Algorithm Engineering. Bridging the Gap Between Algorithm Theory and Practice, Springer Verlag	(CO)

Journal accepted

	Year	Authors	Title	Venue	
JA1	-	M.A. Abam, P.K. Agarwal, M. de Berg and H. Yu	Out-of-order event processing in kinetic data structures	To appear in Algorithmica	(PR)(CO)
JA2	-	M.A. Abam and M. de Berg	Kinetic spanners in R^d	To appear in Discrete & Computational Geometry	(PR)(CO)
JA3	-	G. S. Brodal, B. Gfeller, A.G. Jørgensen and P. Sanders	Towards Optimal Range Median	To appear in Theoretical Computer Science	(PR)(CO)
JA4	-	M. Kutz, G.S. Brodal, K. Kaligosi and I. Katriel	Faster Algorithms for Computing Longest Common Increasing Subsequences	To appear in Journal of Discrete Algorithms	(PR)(CO)
JA5	-	M.A. Bender, G.S. Brodal, R. Fagerberg, D. Ge, S. He, H. Hu, J. Iacono and A. López-Ortiz	The Cost of Cache-Oblivious Searching	To appear in Algorithmica	(PR)(CO)
JA6	-	P.K. Agarwal, L. Arge and K. Yi	I/O-Efficient Batched Union-Find and Its Applications to Terrain Analysis	To appear in ACM Transactions on Algorithms	(PR)(CO)
JA7	-	P. Afshani, C. Hamilton and N. Zeh	Cache-Oblivious Range Reporting With Optimal Queries Requires Superlinear Space	To appear in Discrete & Computational Geometry	(PR)(CO)
JA8	-	P. Afshani, C. Hamilton and N. Zeh	A Unified Approach for Cache-Oblivious Range Reporting and Approxi-mate Range Counting	To appear in Computational geometry: Theory and applications	(PR)(CO)
JA9	-	E.D. Demaine, M. Hajiaghayi and B. Mohar	Approximation Algorithms via Con-traction Decomposition	To appear in Combinatorica	(PR)(CO)
JA10	-	E.D. Demaine, M. Hajiaghayi, H. Mahini and M. Zadimoghaddam	The Price of Anarchy in Network Creation Games	To appear in ACM Transactions on Algorithms	(PR)(CO)
JA11	-	J. Cardinal, E.D. Demaine, S. Fiorini, G. Joret, S. Langerman, I. Newman and O. Weimann	The Stackelberg Minimum Spanning Tree Game	To appear in Algorithmica	(PR)(CO)
JA12	-	E.D. Demaine, M.L. Demaine, J. Iacono and S. Langerman	Wrapping Spheres with Flat Paper	To appear in Computational Geometry: Theory and Applications	(PR)(CO)
JA13	-	G. Aloupis, S. Collette, M. Damian, E.D. Demaine, R. Flatland, S. Langerman, J. O'Rourke, S. Ramaswami, V. Sacristan and S. Wuhrer	Linear Reconfiguration of Cube-Style Modular Robots	To appear in Computational Geometry: Theory and Applications	(PR)(CO)
JA14	-	E.D. Demaine, F. Gomez-Martin, H. Meijer, D. Rappaport, P. Taslakian, G.T. Toussaint, T. Winograd and D.R. Wood	The Distance Geometry of Music	To appear in Computational Geometry: Theory and Applications	(PR)(CO)
JA15	-	N. Alon, M. Badoiu, E.D. Demaine, M. Farach-Colton, M. Hajiaghayi and A. Sidiropoulos	Ordinal Embeddings of Minimum Relaxation: General Properties, Trees, and Ultrametrics	To appear in ACM Transactions on Algorithms	(PR)(CO)
JA16	-	M.A. Bender, G.S. Brodal, R. Fagerberg, R. Jacob and E. Vicari	Optimal Sparse Matrix Dense Vector Multiplication in the I/O-Model	To appear in Theory of Computing Systems	(PR)(CO)

Journal submitted

	Year	Authors	Title	Venue	
JS1	-	C. Demetrescu, B. Escoffier, G. Moruz and A. Ribichini	Adapting Parallel Algorithms to the W-Stream Model, with Applications to Graph Problems	Submitted to Theoretical Computer Science	-
JS2	-	J. Freixas, X. Molinero, M. Olsen and M. J. Serna	On the Complexity of Problems on Simple Games	Submitted to Applied Mathematics and Computation	-
JS3	-	J. Katajainen and S. S. Rao	A compact data structure for representing a dynamic multiset	Submitted to Information Processing Letters	-
JS4	-	J.E. Moeslund, L. Arge, P.K.Bøcher, B. Nygaard and J-C. Svenning	Remote-sensing-based geographically comprehensive assessment of vegetation-elevation relationships in salt meadows across Denmark and implications for sea-level rise impacts	Submitted to Journal of Applied Ecology	-
JS5	-	P. Afshani and T.M. Chan	Optimal Halfspace Range Reporting in Three Dimensions	Submitted to Discrete & Computational Geometry	-
JS6	-	L. Arge, G.S. Brodal, S. S. Rao	External memory planar point location with logarithmic updates	Submitted to Algorithmica	-

Conference submitted

	Year	Authors	Title	Venue	
CS1	-	F. Gieseke, G. Moruz and J. Vahrenhold	Resilient kd-trees: K-means in space revisited	Submitted to 2010 Conference on Data Mining (ICDM)	-
CS2	-	G. S. Brodal and J. Truelsen	An Implicit Dictionary with the Working Set Property	Submitted to 2010 International Symposium on Algorithms and Computation (ISAAC)	-
CS3	-	D. Ajwani and N. Zeh	Topological Sorting of Large DAGs with Small Path Cover	Submitted to 2010 International Symposium on Algorithms and Computation (ISAAC)	-
CS4	-	S. Daneshpajouh, M.A. Abam, L. Deleuran and M. Ghodsi	Computing Homotopic Simplification in a Plane	Submitted to 2010 International Symposium on Algorithms and Computation (ISAAC)	-
CS5	-	L. Arge, K. D. Larsen, T. Mølhave and F. van Walderveen	Cleaning Massive Sonar Point Clouds	Submitted to 2010 International Conference on Advances in Geographic Information System (ACM-GIS)	-
CS6	-	D.M. Kane and J. Nelson	A Derandomized Sparse Johnson-Lindenstrauss Transform	Submitted to 2011 Symposium on Discrete Algorithms (SODA)	-
CS7	-	G. Moruz and A. Negoescu	Fast strongly competitive randomized paging algorithms	Submitted to 2011 Symposium on Discrete Algorithms (SODA)	-
CS8	-	A.G. Jorgensen, M. Löffler and J. Phillips	Geometric Computations on Indecisive Points	Submitted to 2011 Symposium on Discrete Algorithms (SODA)	-

CS9	-	A.G. Jørgensen and K. D. Larsen	Range Selection and Median: Tight Cell Probe Lower Bounds and Adaptive Data Structures	Submitted to 2011 Symposium on Discrete Algorithms (SODA)	-
CS10	-	P. Afshani and N. Zeh	Improved Space Bounds for Cache-Oblivious Range Reporting	Submitted to 2011 Symposium on Discrete Algorithms (SODA)	-
CS11	-	P. Afshani, G.S. Brodal and N. Zeh	Ordered and Unordered Top-K Range Reporting in Large Data Sets	Submitted to 2011 Symposium on Discrete Algorithms (SODA)	-
CS12	-	G. S. Brodal, P. Davoodi and S. S. Rao	Dynamic Bottleneck Edge Queries on Trees	Submitted to 2010 Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS)	-
CS13	-	G.S Brodal, Ss.Sioutas, K. Tsichlas and C. Zaroliagis	D2-Tree: A New Overlay with Deterministic Bounds	Submitted to 2010 International Symposium on Algorithms and Computation (ISAAC)	-

Manuscripts

	Year	**Authors**	**Title**	**Venue**	
M1	-	J. Nelson and D.P. Woodruff	A Near-Optimal Algorithm for L1-Difference	-	-
M2	-	D. Ajwani, A. Beckmann and U. Meyer	I/O-efficient Approximation of Graph Diameter	-	-
M3	-	L. Arge, H. Blunck and A. Hesselund-Jensen	I/O-efficient Nearest Neighbor Algorithms for Grid DEM Quality Evaluation	-	-
M4	-	L. Arge, H. Blunck and A. Hesselund-Jensen	I/O-efficient quadtree construction and leaf neighbor finding	-	-
M5	-	L. Arge and M. Thorup	RAM Efficient External Memory Algorithms	-	-
M6	-	P. Afshani, P.K Agarwal, L.Arge, K.D. Larsen and J.M. Phillips	Approximate Uncertain Skylines	-	-
M7	-	H. Haverkort and F. van Walderveen	Four-Dimensional Hilbert Curves for R-Trees	-	-
M8	-	M.T. Goodrich and N. Sitchinava	Sorting, Searching, and Simulation in the MapReduce Framework	-	-
M9	-	M.A. Abam, L. Arge and L. Deleuran	Efficiently Computing the Canonical Sequence of a Path and its Applications	-	-
M10	-	P. Afshani	On Approximate Simplicial Depth Queries	-	-
M11	-	M. Olsen and T. Viglas	Link Building	-	-

Total number of publications

Category	Peer Rev.	Non-peer Rev.
Number of journal articles (published and accepted)	39	
Number of conference series	131	7
Number of monographs		2
Number of book chapters	2	2
Other		22

Total number of unpublished publications

Number of journal submissions		6
Number of conference submissions		13
Number of manuscripts		11

Note: Due to the theoretical nature of much of the center research and the relatively fast publication rate of the area, it is somewhat hard to determine exactly which of the publications above result from activities before the establishment of the center. In terms of the conferences, publications C2, C5, C10, C11, C12, C13, C17, C18, C20 are published in conferences with submission deadlines before March 1, 2007, and the bulk of the work on these papers was therefore done before the center was established. Publications C1, C3, C4, C6, C8, C9, C14, C19 were submitted in the spring of 2007 but after the center establishment on March 1. In terms of journals, the 2007-2008 papers J1 through J4 are certainly a result of activities before the establishment of the center, and so are a good number of the 2009 papers J5 through J13.

Bibliometric analysis:
At the time of the first annual report, the center was asked to name the 10 most prestigious publication venues in the field. Since conferences play an equally important role as journals in computer science (and are highly refereed), we indicated the below 10 journals and 10 conferences. We also indicated that the ranking should not be taken too seriously.

Estimated 10 most prestigious journals in the field:
- 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

Estimated 10 most prestigious conferences in the field:
- 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 Randomization and Computation (RANDOM)
- Scandinavian Workshop on Algorithm Theory (SWAT)/Workshop on Algorithms and Data Structures (WADS)
- Workshop on Algorithm Engineering and Experiments (ALENEX)

During the first center period, we have published the following number of papers in the listed *journals*:

Journal of the ACM	0
SIAM Journal on Computing	1
ACM Transactions on Algorithms	6
Discrete & Computational Geometry	5
Algorithmica	8
Journal of Computer and System Sciences	0
Computational Geometry: Theory and Applications	6
ACM Journal of Experimental Algorithmics	1
Theoretical Computer Science	1
Journal of Discrete Algorithms	1

The relatively modest amount of center journal papers is a result of the focus on conferences in the computer science area, combined with an often very long journal publication time. Of the listed journals, the *Journal of the ACM* is very broad (covering all of computer science), while the rest of the journals are algorithms journals. The journals *Discrete & Computational Geometry*, *Computational Geometry: Theory and Applications*, *ACM Journal of Experimental Algorithmics* are somewhat more focused (area wise) than the rest. It should be noted that center researchers have also published in the broader *Communications of the ACM*, *Proceedings of the IEEE* and *Proceedings of the US National Academy of Sciences*.

In terms of *conferences*, we believe that one can categorize 6 of the 10 listed conferences as general (or broad) and the rest as specialized algorithms conferences, and also identify the top conferences in the two categories. Using this categorization, the center's conference publication record during the first center period is as follows:

Top general/broad algorithms conferences

ACM-SIAM Symposium on Discrete Algorithms (SODA)	17
IEEE Symposium on Foundations of Computer Science (FOCS)	12
ACM Symposium on Theory of Computing (STOC)	1

Top of rest of general algorithms conferences

European Symposium on Algorithms (ESA)	8
International Colloquium on Automata, Languages, and Programming (ICALP)	9
Scandinavian Workshop on Algorithm Theory (SWAT) Workshop on Algorithms and Data Structures (WADS)	10

Note that in the top of rest category the *Scandinavian Workshop on Algorithm Theory* (SWAT) and *Workshop on Algorithms and Data Structures* (WADS) are clearly lower ranked than *European Symposium on Algorithms* (ESA) and *International Colloquium on Automata, Languages, and Programming* (ICALP). They are "regional" Scandinavian and Canadian conferences. Center researchers have e.g. also published extensively (12 papers) in the similar Asian conference *International Symposium on Algorithms and Computation (ISAAC).*

Top specialized algorithms conferences

Symposium on Computational Geometry (SoCG)	11
Workshop on Algorithm Engineering and Experiments (ALENEX)	1

Top of rest of specialized algorithms conferences

ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)	7
International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX) International Workshop on Randomization and Computation (RANDOM)	2

Note that *European Symposium on Algorithms* (ESA) listed in the general category above has both a theoretical and algorithm engineering track (with separate program committees), and therefore the conference could also have been listed in the top specialized category, on par with, or even a little above, *Workshop on Algorithm Engineering and Experiments* (ALENEX). Note that center researchers have been on both the ALENEX and ESA algorithm engineering track program committees during the center period, which means that a number of papers could not be submitted to those conferences. Thus several center algorithm engineering papers have also been published in the lower ranked conferences *International Workshop on Experimental Algorithms* (WEA) and *International Symposium on Experimental Algorithms* (SEA). Finally, it is also worth noticing that center researchers have also published 6 paper in the top theoretical database conferences *ACM Symposium on Principles of Database Systems* (PODS) and *International Conference on Database Theory* (ICDT).

The *impact of center's publications* is of course hard to measure. One indication of the breadth of the center work is that altogether the center publications have 238 authors, 46 of which are associated with the center. Only 62 of the 235 center publications are authored solely by center researchers. Citations are obviously an important impact indicator, but maybe mostly over a little longer time period than the first center period. Nevertheless, according to Google Scholar (which is probably the most complete index of computer science publications) center publications have been cited 1195 times, distributed with 231 citation for center papers published in 2007 (all papers cited), 632 for papers published in 2008 (38 papers cited), 245 for papers published in 2009 (53 papers cited) and 87 for papers published in 2010 (22 papers cited). However as always, automated citation counts should be taken with a grain of salt.

g) Patents and applications

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 exploitation (please specify)	Partner (s) if any
0	0	1 (Software agreement with SCALGO)	Scalable Algorithmics (SCALGO)		COWI A/S EIVA A/S

Note: Table *only* includes information for AU employees.

h) Other relevant information documenting international standing of center

This appendix contains tables with information about center awards, center AU funding, center AU editorial work, center conference and seminar contributions and center collaborations. This corresponds to the information the foundation requires in the annual reports (and it is given in the annual report format). Note that much of the information is only given for AU employees, and that due to its volume some of the information might be incomplete.

Awards and acknowledgments

Award	Recipient	Amount in DKK
Appointed MIT Associate professor with Tenure, 2007	Demaine	
Appointed MIT Associate professor with Tenure, 2007	Indyk	
Honorary Doctor of Laws degree, Dalhousie University, 2007	Demaine	
Elected member of Royal Danish Academy of Science and Letters, 2008	Arge	
Honorary Doctor of Science at Aarhus University, 2008	Mehlhorn	
International Francqui Chair of Belgium and Francqui Gold Medal, 2008	Demaine	~400.000
Katayanagi Emerging Leadership Prize, CMU and Tokyo University of Technology, 2008	Demaine	~60.000
Université Libre de Bruxelles Gold Medal, 2008	Demaine	
Paper invited to the "Research Highlights" section of the Communications of the ACM, January 2008.	Indyk	
Best student paper award, Symposium on Foundations of Computer Science 2008	Patrascu	
Elected Distinguished Scientist by Association for Computing Machinery (ACM), 2009	Arge	
George Polya Lecturer, Mathematical Association of America, 2010-2012	Demaine	
Appointed Associate Editor of ACM Journal on Experimental Algorithmics	Arge	
Appointed Associate Professor with Exceptional Qualifications (Lektor MSK), 2009	Brodal	
Google European Doctoral Fellowship, 2010-2013	Larsen	~1.100.000
Natural Sciences and Engineering Research Council of Canada (NSERC) postdoctoral fellowship, 2010-2012	Afshani	~500.000
Danish Minister of Research Elite Researcher Award ("Videnskabsministerens EliteForsk-Pris"), 2010	Arge	1.200.000
Best Paper Award, Symposium on Principles of Database Systems, 2010	Nelson	
Charles & Jennifer Johnson MIT Master's of Engineering Award in Computer Science, 2010	Berinde	
Appointed Associate Editor of IEEE Transactions on Signal Processing, 2010	Indyk	
Winner of several categories in the Sortbenchmark, 2010	Beckmann, Meyer	
EATCS Award 2010	Mehlhorn	

AU external funding

Public Danish funds

Funding body	Purpose	Applicant	Activity period	Amount in DKK
Science Research Council	Algorithms for Processing Massive Datasets: Theory and Practice	Arge	2004-2008	~5.800.000
Danish Directorate for Food, Fisheries and Agri Business	Geology-dependent Variations in Transport Processes with respect to Risk Assessment of Phosphorus Loss	DJF, AU (including Arge), and others	2006-2007	~110.000 (AU part)
Danish National Advanced Technology Foundation (Højteknologifonden)	A platform for Galileo based pervasive positioning	AU (including Arge), AAU, Danish agricultural advisory service and several companies	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
Forsknings- og innovationsstyrelsen, PhD kvalitets-fremmemidler	Post Doc support	Arge (through BRICS)	2009	~250.000
Danish Minister of Research	Elite Researcher Award	AU (for Arge)	2010-2012	1.000.000

Private Danish funds

Funding body	Purpose	Applicant	Activity period	Amount in DKK
Aarhus Universitets Forskningsfond	Center for Interdisciplinary Geospatial Informatics Research	Faculty of Natural Sciences, AU (including Arge)	2009-2012	2.500.000
VELUX Visiting Professor Program	Visit of Norbert Zeh	Arge (through Faculty of Natural Sciences, AU)	2009	80.000

International funds

Funding body	Purpose	Applicant	Activity period	Amount in DKK
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
Google	European Doctoral Fellowship	AU (Arge and Larsen)	2010-2013	~1.000.000
Slovenian Research Agency	Processing of Massive Geometric Data	University of Aribor, University of Primorska, AU (Arge and Brodal) and others	2010-2013	~2.000.000

Note: Tables *only* include information for AU employees.

AU editorial work

Journal

Journal	Number of reviews
Journal of Computational Geometry	0
Journal of Experimental Algorithmics	~5
Journal of Graph Algorithms and Applications	~5
Journal of Discrete Algorithms	0
Algorithmica	~15
Electronic Proceedings in Theoretical Computer Science	1
Theoretical Computer science	4
Encyclopedia of Algorithms	~8

Conference program committee

Conference	Number of reviews
ACM Symposium on Advances in Geographic Information Systems, 2007	~10
International Colloquium on Automata, Languages and Programming, 2007	Chair (~150 sub.)
International Symposium on Spatial and Temporal Databases, 2007	~10
International Conference on Data Engineering, 2007	~10
Algorithm Engineering track of European Symposium on Algorithms, 2007	Chair (~50 sub.)
ACM-SIAM Symposium on Discrete Algorithms (SODA), 2007	38
Workshop on Algorithm Engineering and Experiments, 2007	Co-chair (~60 sub.)
Symposium on Theoretical Aspects of Computer Science, 2007	70
Workshop on Algorithms and Data Structures, 2007	18
IEEE International Symposium on Parallel and Distributed Processing with Applications, 2007	5
International Conference on Foundation of Software Technology and Theoretical Computer Science, 2007	~15
International Symposium on Algorithms and Computation, 2008	3
ACM International Symposium on Advances in Geographic Information Systems, 2008	~10
International Symposium on Algorithms and Computation, 2008	~15
Symposium on Combinatorial Pattern Matching, 2008	9
International Workshop on Experimental Algorithms, 2008	~16
Workshop on Massive Data Algorithmics, 2009	48
ACM International Symposium on Advances in Geographic Information Systems, 2009	~7
International Symposium on Spatial and Temporal Databases, 2009	~6
International Computing and Combinatorics Conference, 2009	~17
International Symposium on Experimental Algorithms, 2009	16
ACM Symposium on Parallelism in Algorithms and Architectures, 2009	22
Workshop on Algorithms and Data Structures, 2009	10
Latin American Theoretical Informatics Symposium, 2009	12
Workshop on Massive Data Algorithmics	16
ACM International Symposium on Advances in Geographic Information Systems, 2010	7
International Symposium on Algorithms and Computation, 2010	18
ACM-SIAM Symposium on Discrete Algorithms, 2010	47
European Symposium on Algorithms, 2010	13
International Workshop on Combinatorial Algorithms, 2010	7
Scandinavian Workshop on Algorithm Theory, 2010	15

Note:

- Tables *only* include information for AU employees.
- In addition to the above, AU employees did approximately 15 other but non-editor journal reviews and approximately 100 non-program committee conference reviews.
- The two senior AU employees also did more than 10 international grant and tenure/career assessments, served on 3 conference steering committees and on one project advisory board, as well as on 7 PhD evaluation committees.

Center conference and seminar contributions

2007

Event	Participant(s)	Contribution	Invited talk
European Symposium of Algorithms	Mølhave	Talk	
Workshop on Algorithms and Data Structures	Mølhave	Talk	
International Symposium on Mathematical Foundations of Computer Science	Jørgensen, Brodal	Talks	
Bertinoro Workshop on Algorithms and Data Structures	Brodal	Talk	X
Oberwolfach meeting on Algorithm Enginnering	Brodal	Talk	X
Conference on Computability in Europe, CiE 2007	Olsen	Talk	
Symposium on Mathematical Foundations of Computer Science	Moruz	Talk	
Oberwolfach Seminar on Algorithm Engineering	Moruz, Meyer	Talks	X
International Conference on Algorithmic Learning Theory	Chang	Talk	
Symposium on Mathematical Foundations of Computer Science	Mehlhorn	Talk	X
Seminar, Microsoft Research, Mountain View	Indyk	Talk	X
European Workshop on Computational Geometry	Demaine	Talk	X
Kyoto International Conference on Computational Geometry and Graph Theory	Demaine	Talk	X
Workshop on Algorithms and Data Structures	Demaine	Talk	
Canadian Conference on Computational Geometry	Demaine	Talk	
IEEE Symposium on Foundations of Computer Science	Patrascu	Talk	
ACM Symposium on Theory of Computing	Patrascu	Talks	
ACM Symposium on Computational Geometry	Demaine, Patrascu	Talks	
International Colloquium on Automata, Languages and Programming	Weimann	Talk	

2008

Event	Participant(s)	Contribution	Invited talk
European Symposium of Algorithms	Mølhave, Arge, Mehlhorn	Talk, PC	
Dagstuhl seminar on Data Structures	Arge, Abam	Talks, PC	
Korean workshop on Computational geometry	Arge	Talk	X
International Symposium on Algorithms and Computation	Arge	PC	
Seminar, University of Koper, Slovenia	Arge	Talk	X
Seminar, Computer University of Copenhagen	Arge	Talk	X
Symposium on Computational Geometry	Rao, Abam, Sidiropoulos, Demaine	Talks	
Scandinavian Workshop on Algorithm Theory	Rao, Demaine, Meyer	Talks	
International Symposium on Algorithms and Computation	Grønlund Jørgensen, Arge	Talk, PC	
Northwest European Programming Contest 2008	Truelsen, Greve, Dalgaard Larsen, Mølhave, Ehlers Nyholm Thomsen	Participation	
International Computing and Combinatorics Conference	Olsen	Talk	
International Workshop on Experimental Algorithms	Ajwani	Talk	
Canadian Conference in Computational Geometry	Abam	Talk	
Seminar, Paderborn University	Abam	Talk	
Seminar, Carleton University	Abam	Talk	
Seminar, Institute of Business and Technology, Aarhus University	Olsen	Talk	
Seminar, NICTA Sydney	Olsen	Talk	
Dagstuhl Seminar on Sublinear Algorithms	Meyer	Talk	X
Seminar of the German Research Cluster Algorithm Engineering	Meyer	Talk	X
Algorithm Engineering Workshop @ Google	Meyer, Moruz	Talks	
Symposium on Theoretical Aspects of Computer Science	Meyer	Talk	
Matching under Preferences, Satellite Workshop of ICALP	Mehlhorn	Talk	X

Kyoto RIMS Workshop on Computational Geometry and Discrete Mathematics	Mehlhorn	Talk	X
The Design and the Elastic Mind Symposium	Demaine	Talk	X
International Workshop on Parameterized and Exact Computation	Demaine	Talk	X
7th International Conference on Information Processing in Sensor Networks	Demaine	Paper	
IEEE Conference on Computational Complexity	Demaine	Paper	
MathFest 2008	Demaine	Talks	X
International Workshop on Approximation Algorithms for Combinatorial Optimization Problems	Demaine	Paper	
Advances in Architectural Geometry	Demaine	Paper	
DARPA InfoChemistry meeting	Demaine	Talk	X
International Workshop on the Algorithmic Foundations of Robotics	Demaine	Papers	
Symposium on Combinatorial Pattern Matching	Demaine	Talk	
IEEE Symposium on Foundations of Computer Science	Nelson, Sidiropoulos, Pătraşcu	Talks	
ACM/SIAM Symposium on Discrete Algorithms	Pătraşcu	Talk	
Seminar, Stanford University	Indyk	Talk	X
Seminar, University of Illinois at Urbana-Champaign	Indyk	Talk	X
Seminar, University of Michigan at Ann Arbor	Indyk	Talk	X
Tutorial, Princeton University	Indyk	Talk	X
Allerton Conference	Indyk	Talk	X
Seminar, State University of New York at Buffalo	Indyk	Talk	X
Seminar, Dartmouth College	Indyk	Talk	X
Summer School, Kent State University	Indyk	Talk	X

2009

Event	**Participant(s)**	**Contribution**	**Invited talk**
Beyond Kyoto: Addressing the Challenges of Climate Change. Science meets Industry, Policy and Public	Moeslund, Arge	Poster	
Mapping and Monitoring of Nordic Vegetation and Landscapes	Moeslund	Poster	
International Symposium on Experimental Algorithms	Hachenberger, Beckmann	Talks	
Symposium on Foundations of Computer Science	Larsen, Afshani, Arge, Do Ba	Talks	
Algorithms and Data Structures Symposium	Mølhave, Demaine	Talks	
Annual Symposium on Computational Geometry	Arge, Brodal, Indyk, Demaine, Mehlhorn, Afshani, Abam (and others)	Talks, Organizing	
Current Trends in Algorithms, Complexity Theory, and Cryptography	Afshani, Arge, Brodal	Talks	X
Workshop on Theory and Many-Cores	Sitchinava	Talk	
Seminar, Cambridge University	Sitchinava	Talk	X
International Symposium on Algorithms and Computation	Revsbæk, Arge, Greve, Jørgensen, Tsakalidis, Negoescu, Moruz, Demaine	Talks	
European Workshop on Computational Geometry	van Walderveen	Paper	
European Symposium of Algorithms	Arge, Brodal, Abam, Mehlhorn, Meyer, Demaine	Steering Committee, Talks	X
International Symposium on Spatial and Temporal Databases	Arge	Talk, PC	X
Symposium on Principles of Database Systems	Arge	Talk	X
Seminar, Duke University	Afshani	Talk	X
Seminar, University of Southern Denmark	Greve	Talk	X
Seminar, IT University Copenhagen	Greve	Talk	X

Workshop on Massive Data Algorithmics	Arge, Brodal, Jørgensen, Ajwani, Beckmann (and others)	Talks, Organizing	
Dagstuhl seminar on Computational Geoemtry	Abam, Mehlhorn, Arge	Talks	X
Dagstuhl seminar on Geometric Networks, Metric Space Embeddings and Spatial Data Mining	Abam	Talk	X
Seminar, Institute of Mathematical Sciences, Chennai	Rao	Talk	X
Dagstuhl seminar on Graph Search Engineering	Ajwani, Meyer	Talks	X
Workshop on Algorithm Engineering and Experiments	Meyer	Paper	
International Symposium on Parallel and Distributed Processing	Beckmann	Paper	
International Conference on Symbolic and Algebraic Computation	Mehlhorn	Paper	
Seminar of the German Research Cluster Algorithm Engineering	Moruz	Talk	X
Kurt Mehlhorn 60th Birthday Colloquium	Meyer	Talk	X
Colloquium on Automata, Languages and Programming	Mehlhorn, Demaine, Nelson	Talks	X
Seminar, University of Dortmund	Moruz	Talk	X
Information Theory and Applications Workshop	Indyk	Talk	X
Seminar, Rice University	Indyk	Talk	X
EPFL Summer Research Institute	Indyk	Talk	X
Conference on Learning Theory	Indyk	Talk	X
Workshop on Algorithms for Data Streams	Indyk, Nelson	Talks	X
New Directions in Applied Mathematics Workshop	Indyk	Talk	X
Microsoft Research-India Theory Day	Indyk	Talk	X
Complexity reading group, Microsoft Research	Nelson	Talk	X
Seminar, University of Michigan	Nelson	Talk	X
China Theory Week	Nelson	Talk	X
Seminar, Cal. Tech.	Nelson	Talk	X
Seminar, UCLA	Nelson	Talk	X
Seminar, IBM Almaden	Nelson	Talk	X
Seminar, IT University of Copenhagen	Nelson	Talk	X
DIMACS/DyDan Workshop on Streaming, Coding, and Compressive Sensing	Nelson	Talk	X
Conference of BRIDGES: Mathematical Connections in Art, Music, and Science	Demaine	Talk	X
Workshop on Algorithmic Aspects of Wireless Sensor Networks	Demaine	Talk	X
Japan Conference on Computational Geometry and Graphs	Demaine	Talks, Papers	X
British Science Festival	Demaine	Talk	X
British Origami Society Autumn Convention	Demaine	Talk	X
International Fab Lab Forum and Symposium on Digital Fabrication	Demaine	Talk	X
Annual Meeting of the American Association for Advancement of Science	Demaine	Talk	X
Workshop on Internet & Network Economics	Demaine	Paper	
IEEE/RSJ International Conference on Intelligent Robots and Systems	Demaine	Paper	
European Conference on Mobile Robots	Demaine	Paper	
Canadian Conference on Computational Geometry	Demaine	Papers	
European Workshop on Computational Geometry	Demaine	Paper	
International Symposium on Theoretical Aspects of Computer Science	Demaine	Talk, Papers	
ACM-SIAM Symposium on Discrete Algorithms	Demaine	Talk, Papers	
Seminars, Vrije University, Universite Catholique de Louvain, Universite Libre de Bruxelles	Demaine	Talk	X
Seminar, Northwestern University	Demaine	Talk	X

2010

Event	Participant(s)	Contribution	Invited talk
European Symposium of Algorithms	Sitchinava, Arge, Ajwani, Davoodi, Ajwani, Meyer	Talks, Steering committee	
International Parallel & Distributed Processing Symposium	Sitchinava	Talk	
Workshop on Massive Data Algorithmics	Sitchinava, Larsen, Arge, Revsbæk, Brodal, Meyer, Negoescu, Nelson	Talks, Organization	
Dagstuhl seminar on Data Structures	Larsen, Arge, Brodal, Mehlhorn, Meyer, Demaine	Talks, Organization	X
Symposium on Computational Geometry	Larsen, Arge, Revsbæk	Talks	
International Colloquium on Automata, Languages and Programming	Larsen	Talk	
International Conference on Computing for Geospatial Research & Application	Arge, Revsbæk	Talk	
Workshop on Optimal Data Structures for Efficient Organization and Retrieval of Massive Spatial Data	Afshani, Tsakalidis	Talks	X
Symposium on Discrete Algorithms	Arge, Brodal, Nelson, Price, Demaine, Nelson	PC, Paper	
Seminar, UC Irvine	Sitchinava	Talk	X
International Conference on Database Theory	Tsakalidis	Talk	
Seminar, Carleton University	Davoodi	Talk	X
Seminar, University of Leicester	Davoodi	Talk	X
Seminar, University of Waterloo	Tsakalidis	Talk	X
Scandinavian Workshop on Algorithm Theory	Meyer, Zadimoghaddam	PC	
Symposium on Experimental Algorithms	Meyer	PC	
Workshop on Algorithm Engineering and Experiments	Meyer	PC	
Conf. on Energy-Efficient Computing and Networking	Beckmann	Poster	
Green Computing Conference	Beckmann	Paper	
Seminar of the German Research Cluster Algorithm Engineering	Meyer	Talk	X
Chinese-German Workshop on Algorithm Engineering	Mehlhorn, Meyer	Talks	X
Theory Day, Open University, Tel Aviv	Indyk	Talk	X
Latin American Theoretical Informatics Symposium	Indyk	Talk	X
Workshop on Modern Massive Data Sets	Indyk	Tutorial	X
Seminar, Bonn University	Indyk	Talk	X
Symposium on Theory of Computing	Indyk	PC	
Seminar, Hebrew University	Nelson	Talk	X
Seminar, Technion	Nelson	Talk	X
Seminar, University of Maryland	Nelson	Talk	X
Conference on Principles of Database Systems	Nelson	Talks	
Bellairs Winter Workshop on Computational Geometry	Demaine	Organizer	
The Entertainment Gathering	Demaine	Talk	X
Joint Mathematics Meetings	Demaine	Talks	X
Oberwolfach Graph Theory Meeting	Demaine	Talk	X
British Colloquium for Theoretical Computer Science	Demaine	Talk	X
Gathering for Gardner	Demaine	Talk	X
Annual Hudson River Undergraduate Mathematics Conference	Demaine	Talk	X
International Workshop on Graph-Theoretic Concepts in Comp. Science	Demaine	Talk	X
Symposium on Parallelism in Algorithms and Architectures	Zadimoghaddam	Talks	
International Conference on Origami in Science, Mathematics and Education	Demaine, Zadimoghaddam	Papers, talks	X
Symposium on Foundations of Computer Science	Nelson	Talk	

Note: Tables *only* contain information of events where there was center contribution, and only events *outside* the center sites. Center employees attended many more events than the ones listed.

Collaboration

Collaborator	Title	Output/ results	Danish university/ institution	Foreign university/ institution	Danish company	Foreign company
Cofman.com, Denmark	Internet Search Engines				X	
COWI A/S and Peder Klith Bøcher (Faculty of Agricultural Science, AU), Denmark	Efficient Handling of Massive Heterogeneous Terrain Data	Grant + Efficient terrain processing algorithms and TerraStream software modules	X		X	
BNR A/S, Denmark	GIS in traffic management	Software			X	
Eiva A/S, Denmark	Index construction and Sonar data cleaning	Publication, software			X	
Scalable Algorithmics ApS (SCALGO), Denmark	I/O-efficient terrain algorithms and software				X	
DHI, Denmark	Effektiv udnyttelse af massive sensor data i hydrologiske modeller (Samfinancieret PhD stipendie)	Grant			X	
University of Bergen, Norway, Lund University, Linköping University, KTH Stockholm, Chalmers University of Technology, Goteborg University, Sweden, Aarhus University, University of Southern Denmark, Denmark, IT University of Copenhagen, University of Iceland, Iceland, University of Helsinki, Finland, Vilnius University, Lithuania, University of Latvia, Latvia, Steklov Institute of Mathematics at St.Petersburg, Russia	Nordic Algorithms Network (supported by NordForsk)	Grant	X	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	Project		X		X
Jørgen E. Olesen et al. (Faculty of Agricultural Sciences), Erik Jeppesen et al. (National Environmental Research Institute, AU), Nielse Peter Revsbech (Institute of Biological Sciences, AU), Denmark	iCLIMA (integrated Climate-Ecosystem research: Impact and Adaptation). AU UNIK proposal.	Proposal	X			

Andrej Brodnik (Primorska), Solvenia	Processing of massive geometric data	Grant		X		
Several EU partners	DELIS: Dynamically Evolving Large-Scale Information Systems	Grant		X		
German Algorithm Engineering Cluster	Selected Topics in Algorithm Engineering	Grant		X		
Uri Zwick (Tel Aviv), Israel	Graph Algorithms: Theory and Practice	Grant		X		
Lufthansa Systems, Frankfurt, Germany	Efficient shortest-paths computations with dynamic weights	Proposal				X
GSI Helmholtz Centre for Heavy Ion Research, Darmstadt, Germany	Foundations of memory-efficient information processing for FAIR computing	Proposal		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"	Grant + publications	X		X	
S. Muthu Muthukrishnan (Google and Rutgers University), USA	Identification of data analysis collaboration projects between MADALGO and MassDAL			X		X
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")	Publications	X			
Pankaj K. Agarwal and Bardia Sadri (Duke), Andrew Danner (Swarthmore College), H. Mitasova (NC State University), USA, Ke Yi (HKUST), China	I/O-efficient terrain algorithms and TerraStream	Publications, as well as the TerraStream software package		X		
Jan Vahrenhold (TU Dortmund), Germany and Andrew Danner (Swarthmore College), USA	TPIE	TPIE software package		X		
Martin Smerek (Masaryk University), Czech Republic	I/O-efficient software for BDD manipulation	Software		X		

Rajeev Raman (University of Leicester), Great Britain, Alexander Golynski (Google), USA, Roberto Grossi, Alessio Orlandi (University of Pisa), Italy,Arash Farzan (Waterloo), Canada, Srinivasa S. Rao (Seoul National University),South Korea	Succinct data structures	Publications		X		
Giuseppe F. Italiano et al. (Rome), Italy, Rolf Fagerberg (University of Southern Denmark), Denmark, Jan Vahrenhold, Fabian Gierseke (TU Dortmund), Germany	Resilient data structures	Publications		X		
Riko Jacob (TU Munich), Germany, Norbert Zeh (Dalhousie University), Canada, Itay Malinger, Sivan Toledo (Tel-Aviv University), Israel	Flash Memory Models	Publications		X		
Norbert Zeh (Dalhausie), Canada, Riko Jacob (Munich), Knut Reinert (FU Berlin), Rolf Backofen (Freiburg), Germany, Rolf Fagerberg (University of Southern Denmark), Rasmus Pagh (IT University Copenhagen), Denmark, Michael Bender (Stony Brook), USA	I/O-efficient algorithms	Publications	X	X		
Norbert Zeh (Dalhausie), Ian Munro (Waterloo), Canada, John Iacono (NYU), USA, Stefan Lagnerman (Bruxelles), Belgium	Cache-Oblivious algorithms	Publications		X		
Anna Gilbert, Martin Strauss (U Michigan), David Woodruff (IBM Almaden), Andrew McGregor (UCSD), Graham Cormode (AT&T), Rich Baraniuk (Rice), USA, Milan Ruzic (ITU), Denmark	Streaming algorithms, sparse recovery, compressed sensing	Publications, software libraries		X		X
Mike Goodrich (UC Irvine), USA, Norbert Zeh (Dalhausie), Canada	Algorithms for private-cache chip multiprocessors	Publications		X		
Mark de Berg and Herman Haverkort (TU Eindhoven), Netherlands, Timothy Chan (Waterloo), Canada, and Jeremy Barbay (Universidad de Chile), Chile	Geometric algorithms and data structures	Publications		X		

Joachim Gudmundsson (NICTA), Taso Viglas (University of Sydney), Australia, Loukas Georgiadis (HP Lab), USA	Graph algorithms	Publications		X		X
Mark de Berg (TU Eindhoven) The Netherlands, Mohammad Farshi, Michiel Smid (Carleton), Canada, Joachim Gudmundsson (NICTA), Australia	Geometric spanners	Publications		X		
Hamish Carr (Leeds), England	Scalar field simplification			X		
Hee-Kap Ahn (Pohang University of Science and Technology), Sang Won Bae (KAIST) South Korea and Marc Scherfenberg (Freie Universität, Berlin) Germany, Maarten Löffler (UC Irvine), Jeff M. Phillips (Utah), USA	Geometric Computations on imprecise Points	Publications		X		
Thomas Mailund, Christian Nørgaard Storm Pedersen (Bioinformatics Research Center, Aarhus University), Denmark	Algorithms for Evolutionary Trees	Publications	X			
Shervin Daneshpajouh, Alireza Zarei and Mohammad Ghodsi (Sharif), Iran, Mohammad Ali Abam (Dortmund), Germany, Mark de Berg (TU Eindhoven), Netherlands	Line simplification	Publications		X		
Rajeev Raman (University of Leicester), UK, Beat Gfeller (ETH Zurich), Switzerland, Peter Sanders (Karlsruhe Institute of Technology), Germany, Srinivasa Rao (Seoul), South Korea	Range Minimum and Median Data Structures	Publications		X		
Kostas Tsichlas, Apostolos N. Papadopoulos (Thessaloniki), Spyros Sioutas (Ionian), Alexis Kaporis (Patras), Greece	Probabilistic 3-Sided Range searching	Publications		X		
Rolf Fagerberg (University of Southern Denmark), Denmark, Alex López-Ortiz (Waterloo), Canada	Sorted Range Reporting	Publication	X	X		
Taso Viglas (Sydney), Australia	Link Building	Manuscript		X		
Xavier Freixas, Xavier Molinero Albareda, Maria Serna (Barcelona), Spain	Simple Games	Publications		X		
Mohammad Farshi, Michiel Smid (Carleton), Canada, Paz Carmi (Ben-Gurion), Israel	Semi-separated Pair Decompositions	Publication		X		

Mark de Berg (TU Eindhoven), The Netherlands, Sheung-Hung Poon (Tsing Hua University), Taiwan	Fault-Tolerant Conflict-Free Coloring	Publication		X		
Jyrki Katajainen (University of Copenhagen), Denmark	Dynamic Dictionaries for Multisets	Publication	X			
Lars Engebretsen, Google Zurich	Graph Clustering					X
Henning Meyerhenke (Paderborn), Germany	Realistic Computer Models	Publication		X		
Group of Peter Sanders (Karlsruhe), Germany	STXXL, algorithm engineering and energy-efficient sorting	Publications + software		X		
Robert Krauthgamer (Weizman), Israel, Jiri Matousek (Prague), Chech Republich	Embeddings	Publications		X		
Ilias Diakonikolas (Columbia), USA	Pseudorandomness	Publication		X		
Atri Rudra (SUNY), USA	Group testing	Publication		X		
Mihai Badoiu (Google), Martin Farach-Colton (Rutgers) and MohammadTaghi Hajiaghayi (AT&T Research), USA, Anastasios Sidiropoulos (Toronto), Canada, Noga Alon (Tel Aviv), Israel, Morteza Zadimoghaddam (Sharif University), Iran	Ordinal Embeddings	Publication		X		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 and Takehiro Ito (U. Tokyo), Ryuhei Uehara (JAIST), Yushi Uno (Osaka), Japan, Sandor Fekete (TU Braunschweig), Germany, Nicholas J. A. Harvey (Waterloo), Canada	Folding	Publications		X		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

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		
Joseph O'Rourke (Smith College), USA	Computational Geometry	Publications and book		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, Morteza Zadimoghaddam and Shayan Oveisgharan (Sharif U. Technology), Iran, Daniel Marx, Noga Alon (Tel Aviv U.), Israel, Iran, Robert Hearn (Dartmouth), Tom Leighton (Akamai), USA, Bojan Mohar (British Columbia), Canada	Game theory and approximation algorithms	Publications		X		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		

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		
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		
Greg Aloupis (McGill U.), Jean Cardinal (U. Libre Bruxelles), Sébastien Collette (U. Libre Bruxelles), Martin L. Demaine (MIT), Muriel Dulieu (Polytechnic U.), Ruy Fabila-Monroy (U. National Autonoma Mexico), Vi Hart (vihart.com), Ferran Hurtado (U. Polytecnica Catalunya), Stefan Langerman (U. Libre Bruxelles), Maria Saumell (U. Polytecnica Catalunya), Carlos Seara (U. Polytecnica Catalunya), Perouz Taslakian (U. Libre Bruxelles)	Matching Points with Things	Publication		X		X

E. Hawkes (Stanford), B. An (MIT), N. Benbernou (MIT), H. Tanaka (Harvard), S. Kim (MIT), D. Rus (MIT), R. Wood (Harvard), USA	Programmable matter by folding	Publication		X		
Danny Hermelin, Gad M. Landau (Heifa), Shir Landau (Bar-Ilan University), Israel	Accelerating Edit-Distance Computation via Text-Compression	Publications		X		

Note: Some entries in the table correspond to extensive collaborations, whereas others are collaborations on a single paper.

i) Information on center staffs involvement in educational activities

AU undergraduate level courses

Title and date of activity	ECTS	Course length (hours)	Number of participants
BSc course: Introduction to Algorithms 1, Spring 2007	5	25	115
BSc course: Introduction to Algorithms 2, Spring 2007	5	25	115
BSc course: Algorithms and Data Structures 1, Spring 2008	5	28	~130
BSc course: Algorithms and Data Structures 2, Spring 2008	5	28	~130
BSc course: Computer Science in Perspective - 1 out of 7 lecturers, Spring 2008	5	10	~130
BSc course: C++/C programming, Institute of Business and Technology, Spring 2008	5	60	~20
BSc course: Algorithms and Data Structures 1, Spring 2009	5	28	170
BSc course: Algorithms and Data Structures 2. Spring 2009	5	28	135
BSc course: Basic Software Development 1, Institute of Business and Technology, Fall 2009	5	60	15
BSc course: Basic Software Development 2, Institute of Business and Technology, Spring 2009	5	60	15
BSc course: Computer Science in Perspective, Fall 2009	5	10	120
DM Programming Contest Coaching, 2009			20
NWERC Programming Contest Coaching, 2009			6
BSc course: Algorithms and Data Structures 1, Spring 2010	5	25	~165
BSc course: Algorithms and Data Structures 2, Spring 2010	5	25	~135
DM Programming Contest Coaching, 2010			~20
NWERC Programming Contest Coaching, 2010			~6

AU graduate level courses

Title and date of activity	ECTS	Course length (hours)	Number of participants
MSc course: I/O-efficient algorithms, Spring 2007	10	36	19
MSc course: Advanced Data Structures, Fall 2007	10	42	29
MSc course: I/O-efficient algorithms, spring 2008	10	39	17
MSc course: Computational Geometry, Fall 2008	10	42	23
PhD course: Advanced Computational	5	18	6

Geometry, Fall 2008			
MSc course: I/O-efficient algorithms, Spring 2009	10	36	12
PhD course: Advanced Randomized Algorithms, Spring 2009	5	21	10
MSc course: Advanced Data Structures, Fall 2009	10	42	50
PhD Course: I/O-Efficient Graph Algorithms, Spring 2010	5	21	~5
PhD Course: Advanced Range Searching, Spring 2010	5	21	~10

Graduate level classes taught outside center (e.g. summer schools)

Title and date of activity	ECTS	Course length (hours)	Number of participants
I/O-efficient algorithms and data structures, four lectures at NORFA summer school on Algorithmic Data Analysis, Spring 2007		4	73
Efficient nearest neighbor search algorithms, four lectures at NORFA summer school on Algorithmic Data Analysis, Spring 2007		4	73
MADALGO Summer School on Data Stream Algorithms, Summer 2007		~20	70
Algorithms for Advanced Processor Architectures, PhD course, IT University in Copenhagen, Summer 2008	5	24	~30
MADALGO Summer School on Cache-Oblivious Algorithms, Summer 2008		~20	60
Explicit Constructions in High-Dimensional Geometry, Summer School on Probabilistic and Fourier Methods in Geometric Functional Analysis, Kent State University, Summer 2008		4	60
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
Invited tutorial on Worst-Case Efficient Range Search Indexing, ACM Symposium on Principles of Database Systems, Fall 2009		2	~30
MADALGO Summer School on Geometric Data Structures, Summer 2010		~20	~50

Notes: Teaching by center employees at MIT, MPI and FRA is *not* listed.

j) Public outreach

Electronic media

Media	Date	Type	Subject/ Title	Contributor
Version2	November 24, 2006	Interview	Nyt dansk center for de store regnestykker	Arge
jp.dk	December 4, 2006	Article	Forsker i massivt angreb på data-havet	Arge
TV2 østjylland	February 4, 2007	Feature	Århus vil være førende	Arge
TV2 østjylland	February 11, 2007	Feature	På besøg i fremtiden	Arge
JPInternetavisen	August 21, 2007	Article	Regning for viderekomne	Arge
Information	August 24, 2007	Article	Nyt forskningscenter: Forskning i store datamængder	Arge
it-vest NYT	October 2007	Article	Teoretiske dataloger in action	Arge
Ingeniøren	December 5, 2007	Article	Er en supercomputer for dyr? Så brug en superalgoritme	Arge
Christiansborgkontakt	December 6, 2007	Interview	Gå ikke over åen efter data	Arge
videnskab.dk	June 16, 2008	Article	Forskere vil redde os fra data-oversvømmelse	Arge
videnskab.dk	March 11, 2009	Feature	Dansk software skal varsle oversvømmelser	Arge
Computerworld	March 12, 2009	Feature	Ny software holder danskerne tørskoede	Arge
TV2.dk	March 12, 2009	Feature	Ny software holder danskerne tørskoede	Arge
Version2	September 18, 2009	Feature	AU-forskere satte turbo på GIS-beregninger med ram-optimering	Arge
Ingeniøren.dk	September 18, 2009	Feature	Dansk algoritme sætter turbo på digitale terrænmodeller	Arge
videnskab.dk	January 7, 2010	Feature	Hukommelsesfejl skal løses med algoritmer	Brodal
tv2oj.dk	Januar 27, 2010	Feature	Store priser til århusianske eliteforskere	Arge
videnskab.dk	January 27, 2010	Feature	Minister belønner forskere for elite-indsats	Arge
Computerworld	January 27, 2010	Feature	Dansk algoritme-forsker får 1,2 millioner kroner	Arge
cs.au.dk	February 17, 2010	Video	Rollemodeller	Larsen
cs.au.dk	June 28, 2010	Video	Ud af boksen - forskning i algoritmer	Arge
jpaarhus.dk	August 3, 2010	Feature	Google-priser til Aarhus-dataloger	Larsen

Press

Media	Date	Type	Subject/ Title	Contributor
JPÅrhus	December 4, 2006	Article	Forsker i massivt angreb på data-havet	Arge
CAMPUS	No 3/2007	Article	Opskriften på viden	Arge
IT-avisen ComON	Marts 1, 2007	Article	Algoritme-elite samles i dansk center	Arge
IT-avisen ComON	Marts 1, 2007	Article	Behov for algoritmer	Arge
News paper supplement in 6 major danish newspapers	July 9, 2007	Article	Frokost med de store	Arge
JPÅrhus	August 21, 2007	Article	Regning for viderekomne	Arge
Berlingske Tidende	August 24, 2007	Article	Genvejen til hurtigere resultater begynder i Århus	Arge
Nyhedsavisen	August 24, 2007	Article	Løs et problem ved at lave mindst muligt	Arge
Århus Stiftstidende	September 1, 2007	Article	Genvejen til hurtigere resultater begynder i Århus	Arge
Børsen	October 2, 2007	Article	Ny grundforskning skal gøre computerne smartere	Arge
Politiken	October 7, 2007	Names	Forvalter et forskningsbudget på 60 mio. kr.	Arge
Aktuel Naturvidenskab	No 5, 2007	Feature	Gå ikke over åen efter data	Arge

JP	October 7, 2007	Names		Arge
Newletter	November 27, 2008	Feature	Daimi topper i programmeringskonkurrence	Mølhave
Weekendavisen	April 25, 2008	Article	Logaritmen af n	Arge
Jyllandsposten	April 5, 2009	Feature	Ny software varsler oversvømmelser	
Geoforum	May 1, 2009	Article	GIS ved MADALGO	Arge, Mølhave
Ingeniøren	September 18, 2009	Feature	Dansk algoritme sætter turbo på digitale terrænmodeller	Arge
Weekendavisen	October 2, 2009	Feature	Lysindfald ganget med vand	Arge
Aarhus Stiftstidende	December 16, 2009	Front page	Klimaændringer vil drukne Århus	Moeslund
Aarhus Stiftstidende	December 16, 2009	Interview	Ny målemetode tegner et nøjagtigt højdekort	Moeslund
Aarhus Stiftstidende	December 16, 2009	Feature	Klimaændringer slår hårdt i Risskov og Egå	Moeslund
Jyllandsposten	January 28, 2010	Feature	Århus-forskere i top	Arge
Aarhus Stiftstidende	January 28, 2010	Feature	Århusianske forskere vælter sig i priser	Arge
Politiken	April 18, 2010	Feature	Informationssamfundet drukner i … data	Arge
JP Århus	August 3, 2010	Feature	Google priser til Århus dataloger	Larsen

Other

Type	**Date**	**Subject/Title**	**Contributor**
Presentation (teknologisk institut)	June 14, 2007	Effektiv håndtering og analyse af MASSIVE datasæt	Arge
Danske Bank	January 8, 2007	Massive Data Algorithmics	Brodal
UNF Ålborg	February 7, 2007	Internet search engines	Brodal
Gymnasiepraktik	November 2007	Algorithmic Exercises	Brodal
Teknisk Naturvidenskabeligt Gymnasium i Herning	September 2007	Algorithmic Exercises	Brodal
8th grade primary school students	January 2007	Internet search engines	Brodal
Presentation at Geoforum dagseminar om højdedata	April 24, 2008	Algoritmer til håndtering af store datamængder	Arge
Presentation at Nordic Geodetic Commission taskforce workshop	December 11, 2008	Algorithms for Handling Massive (Height) Datasets	Arge
Exhibition: Forskningens Døgn 2008	April 24, 2008	"Oversvømmelse! Hvor? Hvor meget?"	Mølhave, Blunck, Revsbæk
Company lecture, Danske Bank	January 17, 2008	Massive Data Udfordringer	Brodal
Gymnasiepraktik	October 21, 2008	Exercises in Algorithms	Brodal
Presentation at "Kortdage 2009"	November 19, 2009	Hvor løber vandet hen? Oversvømmelsesberegninger på store højdemodeller	Arge
Presentation at "COWI GIS konference 2009"	September 23, 2009	Oversvømmelse og vandstrømning	Arge
Presentation at "Datalogforeningen - faglig aften"	May 27, 2009	MASSIVE Data Algorithmics	Arge
Presentation at "Rotary - Ebeltoft"	April 27, 2009	Hvornår bliver rådhuset oversvømmet - og hvor kommer vandet fra?	Arge
Presentation at "Geoforum højdedataseminar"	March 11, 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
Posters at AU IT-day	June 3, 2010	Various posters describing MADALGO research	
Presentation at "Ungdommens Naturvidenskabelige Forening,	September 29, 2009	Internet Search Engines	Brodal

Ålborg"			
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
Keynote at AU IT-day	June 3, 2010	Massive data algorithmics	Arge
Group exercises	March 15, 2010	Introduction to robotics to Primary School students	Brodal

Note: Tables *only* include information for AU employees.

k) Selected center publications

The below ten papers are selected in an attempt to both cover the breadth of the center's research and to showcase some of the main results discussed in this self evaluation report. This is obviously not a simple task, and many important and interesting papers could not be selected. The number in Appendix f is listed below each paper, along with a few keywords describing the paper.

1. P. Afshani, L. Arge and K.D. Larsen
   Orthogonal Range Reporting in Three and Higher Dimensions
   Proc. Symposium on Foundations of Computer Science (FOCS), 2009
   C80 (I/O and internal memory data structures)

2. G.S. Brodal, E. Demaine, J.T. Fineman, J. Iacono, S. Langerman and J.I. Munro
   Cache-Oblivious Dynamic Dictionaries with Optimal Update/Query Tradeoff
   Proc. Symposium on Discrete Algorithms (SODA), 2010
   C119 (Cache-oblivious data structures)

3. L. Arge, T. Mølhave and N. Zeh
   Cache-Oblivious Red-Blue Line Segment Intersection
   Proc. European Symposium on Algorithm (ESA), 2008
   C34 (Cache-oblivious algorithms)

4. R. Berinde, A. Gilbert, P. Indyk, H. Karloff and M. Strauss
   Combining Geometry and Combinatorics: A Unified Approach to Sparse Signal Recovery
   Proc. Allerton Conference, 2008
   C53 (Streaming algorithms)

5. D.M. Kane, J. Nelson and D.P. Woodruff
   An Optimal Algorithm for the Distinct Elements Problem
   Proc. Symposium on Principles of Database Systems (PODS), 2010
   C112 (Streaming algorithms – best paper award)

6. A. Danner, T. Mølhave, K. Yi, P.K. Agarwal, L. Arge and H. Mitasova
   TerraStream: From Elevation Data to Watershed Hierarchies
   Proc. ACM International Symposium on Advances in Geographical Information Systems (ACM-GIS), 2007
   C7 (I/O algorithm engineering - TerraStream)

7. D. Ajwani, N. Sitchinava and N. Zeh
   Geometric Algorithms for Private-Cache Chip Multiprocessors
   Proc. European Symposium on Algorithms (ESA), 2010
   C129 (Parallel private-cache)

8. A. Beckmann, U. Meyer, P. Sanders and J. Singler
   Energy-Efficient Sorting using Solid State Disks
   Proc. International IEEE Green Computing Conference, 2010
   C116 (Flash, I/O and crosscutting algorithm engineering – sorting record)

9. M. Patrascu
   Succincter
   Proc. Symposium on Foundations of Computer Science (FOCS), 2008
   C23 (Succinct data structures – best student paper award)

10. P. Afshani, L. Arge and K.D. Larsen
    Orthogonal Range Reporting: Query Lower Bounds, Optimal Structures in 3-d, and Higher Dimensional Improvements
    Proc. Symposium on Computational Geometry (SoCG), 2010
    C124 (Internal memory data structures)