



In Memoriam

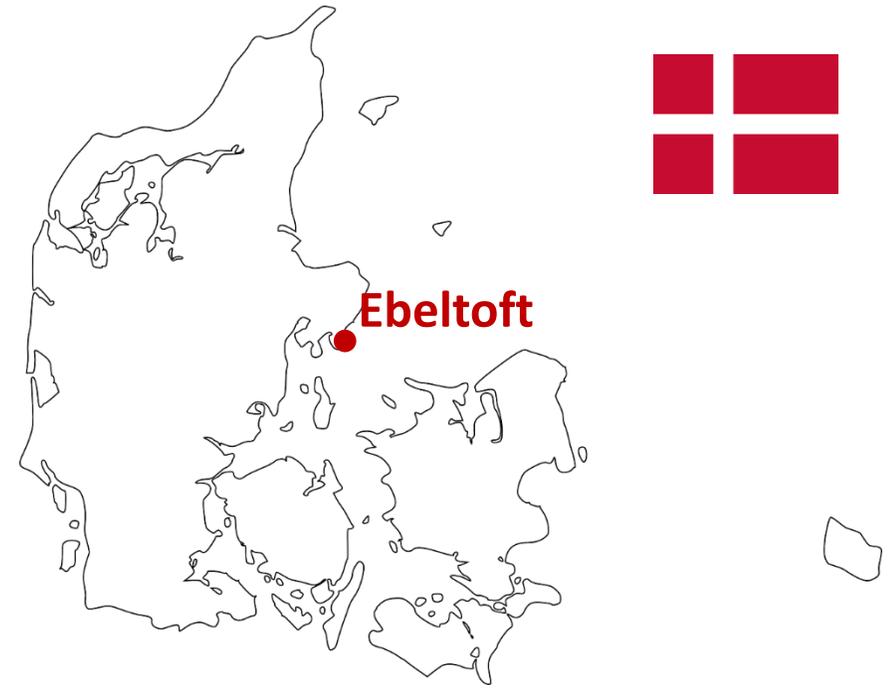
Lars Arge

* 8.10.1967 – † 23.12.2020

Lars Arge Memorial Symposium, October 6, 2021

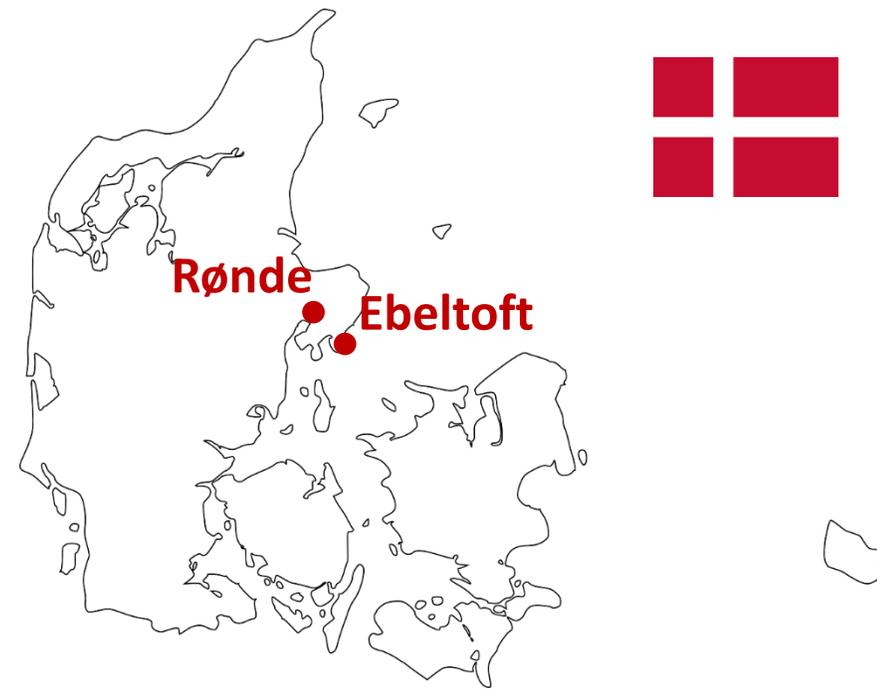
Lars Arge

- Born October 8th 1967 Ebeltoft, Denmark



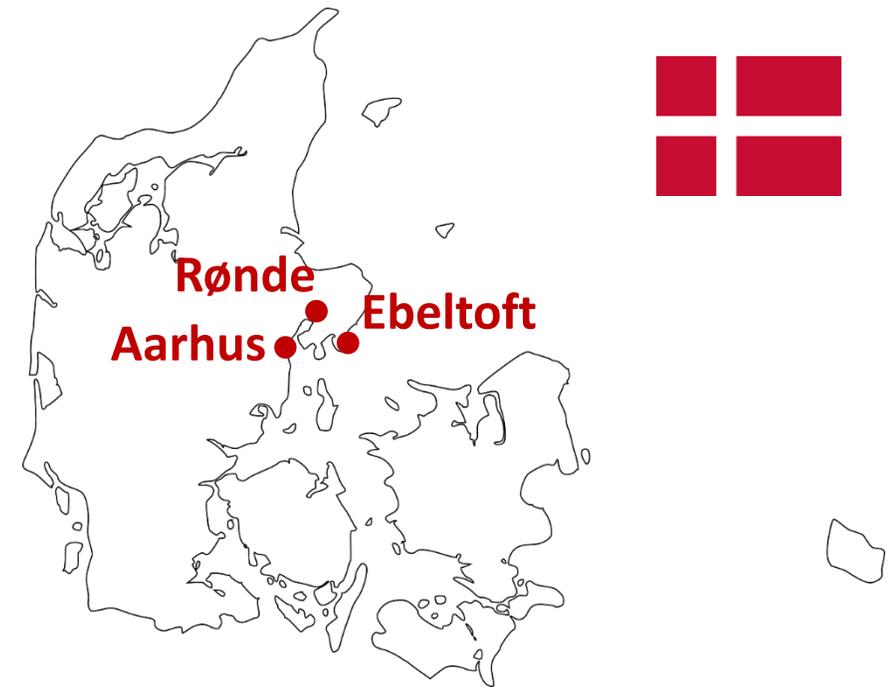
Lars Arge

- Born October 8th 1967 Ebeltoft, Denmark
- High school Rønne



Lars Arge

- Born October 8th 1967 Ebeltoft, Denmark
- High school Rønne
- September 1986 started at Aarhus University



Lars Arge – Aarhus University study

- 1986 started Computer Science and Physics
- 1992 enrolled in the new “4 + 4” PhD program
- 1996 PhD in Computer Science

Chairman of the
Student Council

visited Jeff S. Vitter
8 months at Duke



Lars Arge – PhD thesis Acknowledgments

A single conversation with a wise man is better than ten years of study
Chinese Proverb



- Advisor Erik Meineche Schmidt
 - “He pointed me in the right direction and he entrusted me the freedom...”*
 - “I have had a lot of fun with him drinking beer and discussing politics”*
- Undergraduates that Lars had taught computer science
 - “Not only did they help in “keeping me young”, but it was also a lot of fun.”*
- Mikael Knudsen and Kirsten Larsen
 - “for making the writing of my first paper such a great experience”*
- Sven Skyum
 - “for teaching me computational geometry and for listening to my problems”*
- Jeffrey Scott Vitter and Duke University
 - “The stay at Duke meant everything to me!”*
- Lot of friends, colleagues, secretaries, roommates...

First paper – WADS 1993

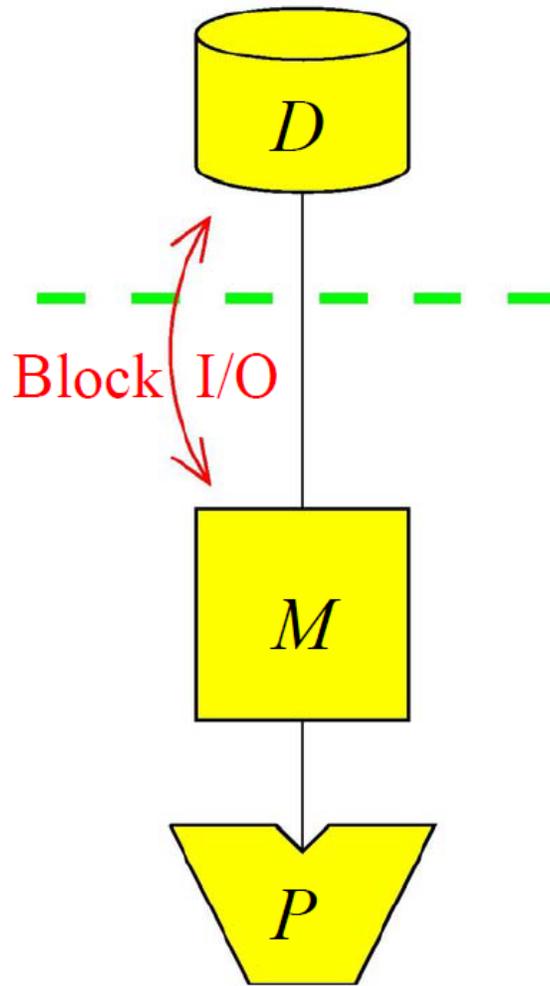
A General Lower Bound on the I/O-Complexity of Comparison-based Algorithms*

Lars Arge, Mikael Knudsen and Kirsten Larsen

Aarhus University, Computer Science Department
Ny Munkegade, DK-8000 Aarhus C.**

Abstract. We show a general relationship between the number of comparisons and the number of I/O-operations needed to solve a given problem. This relationship enables one to show lower bounds on the number of I/O-operations needed to solve a problem whenever a lower bound on the number of comparisons is known. We use the result to show lower bounds on the I/O-complexity on a number of problems where known techniques only give trivial bounds. Among these are the problems of removing duplicates from a multiset, a problem of great importance in e.g. relational data-base systems, and the problem of determining the mode - the most frequently occurring element - of a multiset. We develop algorithms for these problems in order to show that the lower bounds are tight.

I/O-Model



- Parameters

$N = \#$ elements in problem instance

$B = \#$ elements that fits in disk block

$M = \#$ elements that fits in main memory

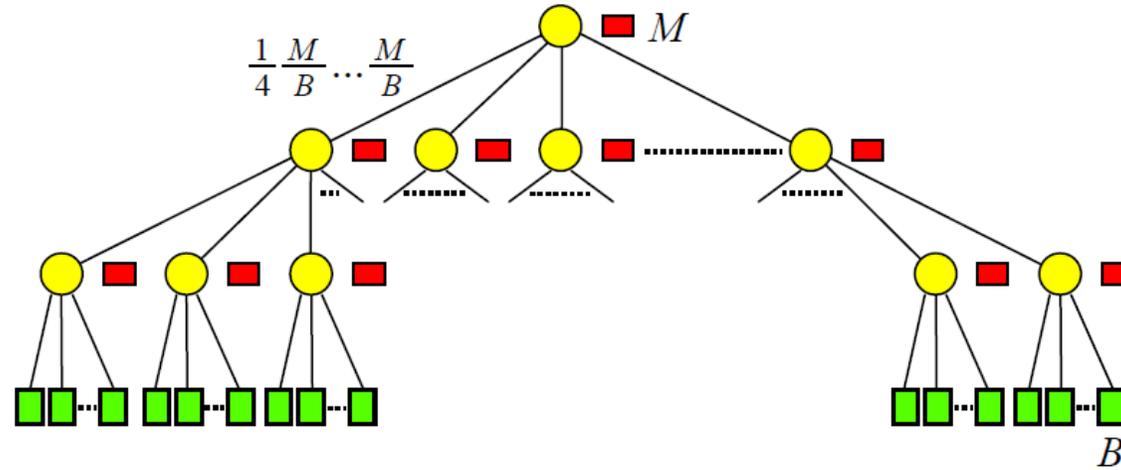
$K = \#$ output size in searching problem

- We often assume that $M > B^2$

- I/O: Movement of block between memory and disk

Basic Buffer-tree

- Definition:
 - B-tree with branching parameter $\frac{M}{B}$ and leaf parameter B
 - Size M buffer in each internal node



- Updates:
 - Add time-stamp to insert/delete element
 - Collect B elements in memory before inserting in root buffer
 - Perform **buffer-emptying** when buffer runs full



Duke University

- 1994 – 1995 Visiting Scholar
- 1996 – 2006 Post Doc / Assistant Prof. / Associate Prof. / Full Prof.
- 2007 – 2015 Adjunct Professor





Returns to Aarhus University

- 2004 – Full professor
- 2004 – 2007 Ole Rømer Scholarship, Danish Research Council
- 2007 – 2017 Center for Massive Data Algorithmics (MADALGO, 70M DKK)
- 2016 – Secretary General, The Royal Danish Academy of Science and Letters



Awards

- 2000 Career Award, US National Science Foundation
- 2008 Elected to the Royal Danish Academy of Sciences and Letters
- 2009 ACM Distinguished Scientist
- 2010 Danish Minister of Research, Elite Research award
- 2011 Elected Member of the Danish Academy of Technical Sciences
- 2012 Fellow of the Association of Computing Machinery (ACM)
- 2015 Knighted by the Queen of Denmark
- 2016 Honorary Doctorate, TU Eindhoven, the Netherlands



2017 Honorary Doctorate Technical University Eindhoven (TU/e)



In the motivation for conferring the degree the Doctorates Board highlights Lars Arges's outstanding achievements in the theoretical foundations of I/O-efficient algorithms, making him a leading expert worldwide, as well as his activities in creating infrastructure for fundamental algorithmic massive-data research and its applications.

Photo: TU Eindhoven / Bart van Overbeeke



The Royal Danish Academy of Sciences and Letters

- Founded 1742
- Appoints the board of the Carlsberg Foundation
- 2008 Lars Arge elected member
- 2015 – Lars Arge member of the Presidium
- 2016 – Lars Arge Secretary General

Contributions

Opened the Academy to the public

Communication of research to the public

Discovered a fraud case and cleaned up



H.C. Andersens Boulevard, Copenhagen

Services

- 2013 – 2019 Member Danish Natural Science Research Council
Vice chair 2015 – 2016, Chair 2016 – 2018
- Editorial
 - Journal editor** ACM Transactions on Spatial Algorithms and Systems, Journal of Experimental Algorithmics, Electronic Proceedings in Theoretical Computer Science, Algorithmica, Journal of Computational Geometry, Journal of Graph Algorithms and Applications, Journal of Discrete Algorithms, Encyclopedia of Algorithms, Theoretical Computer Science, Computational Geometry: Theory and Applications
 - Conference committee (co)-chair** SoCG'15, ICALP'07, ESA'07, SWAT'06, ALENEX'04, Dagstuhl seminars
 - Conference committee member** ALENEX, ACM GIS, CGC, COCOON, CPM, ESA, FAW, FUN, ICALP, ICDE, ICDT, ISAAC, MASSIVE, SoCG, SODA, SPAA, SSTD, STDBM, WADS, WEA, WLDB
 - Steering Committees** SWAT 2004-, ALENEX 2004-2016, ESA 2006-2010 (Chair from 2007), EATCS Council 2012-2019
- Faculty committees
Duke University and Aarhus University
- Boards
Videnskabsklubben, Carlsbergs Mindelegat for Brygger J.C. Jacobsen, Niels Bohr Foundation, Aksel Tovborg Jensen Foundation (chair)

Contribution to education

Graduate students (19)

- Octavian Procopiuc, Laura Toma, Bryan Holland-Minkley, Andrew Danner, Ke (Kevin) Yi, Lars Hvam Petersen, Anders Hessellund Jensen, Thomas Mølhave, Lasse Deleuran, Jesper Moeslund Eshøj, Morten Revsbaek, Kasper Larsen, Freek van Walderveen, Jungwoo Yang, Bryan Wilkinson, Yujin Shin, Mathias Rav, Svend Svendsen, Asger Hautop Drewsen



Post docs (26)

- Norbert Zeh, Herman Haverkort, Loukas Georgiadis, Henrik Blunck, Srinivas Rao, Peter Hachenberger, Mohammad Abam, Deepak Ajwani, Payman Afshani, Nodari Sitchinava, Brody Sandel, Qin Zhang, Elad Verbin, Lap-Kei Lee, Cicimol Alexander, Wei Yu, Constantinos Tsirogiannis, Hossein Jowhari, Darius Sidlauskas, Zhewei Wei, Zengfeng Huang, Wanbin Son, Constantinos Tsirogiannis, Frank Staals, Jesper Sindahl Nielsen, Allan G. Jørgensen

Algorithms courses at Duke (12), Aarhus University (20), outreach talks

A word cloud featuring various computer science terms and names. The most prominent words are 'Proceedings' (large red), 'Arge' (large red), 'Symposium' (large orange), 'Agarwal' (large yellow), 'Data' (large yellow), 'International' (large yellow), 'Algorithms' (large orange), and 'Memory' (large green). Other visible words include 'Computational', 'Structures', 'External', 'Geometry', 'Planar', 'Efficient', 'ACM', 'Vitter', 'I/O-Efficient', 'Workshop', 'European', 'External-Memory', 'Computing', 'Location', 'Discrete', 'Moving', 'Brodal', 'Database', 'Bøcher', 'Vahrenhold', 'Haverkort', 'Grid', 'Systems', 'Svenning', 'FOCS', 'Information', 'Larsen', 'lower', 'Engineering', 'R-Tree', 'Searching', 'Tree', 'Sorting', 'Optimal', 'using', 'SoCG', 'Geographic', 'Priority', 'Scalable', 'IEEE', 'WADS', 'Applications', 'Search', 'ALENEX', 'Schizophrenia', 'Algorithmica', 'ESA', 'Revsbaek', 'Advances', 'Interval', 'special', 'Point', 'R-Trees', 'Queries', 'Zeh', 'Massive', 'External', 'Principles', 'Afshani', 'Journal', 'Geometry', 'Memory', 'Moeslund', 'Yi', 'Management', 'Computation', 'Batched', 'Terrain', 'Danner', 'Computer', 'Dalgaard', 'Chase', 'SODA', 'Tsirogiannis', 'Planar', 'Conference', 'Theory', 'issue', 'Wickremesinghe', 'Erickson', 'Efficient', 'Algorithm', 'Range', 'Cache-Oblivious', 'ACM', 'Moelhave', 'Dynamic', 'Algorithms', 'Yang', 'Vitter', 'Procopiuc', 'Line', 'Indexing', 'ACM-GIS', 'Spatial', 'Graph', 'Berg', 'Experiments', 'I/O-Efficient', 'ACM-SIAM', and 'Goodrich'.



Lars Arge

Professor of Computer Science and MADALGO center director, [Aarhus University](#)

Verified email at cs.au.dk - [Homepage](#)

[Algorithms and data structu...](#)

Cited by

	All
Citations	6956
h-index	46
i10-index	87

[ARTICLES](#) CITED BY

TITLE	CITED BY	YEAR
The influence of Late Quaternary climate-change velocity on species endemism B Sandel, L Arge, B Dalsgaard, RG Davies, KJ Gaston, WJ Sutherland, ... Science 334 (6056), 660-664	554	2011
Indexing moving points PK Agarwal, L Arge, J Erickson Journal of Computer and System Sciences 66 (1), 207-243	429	2003
The priority R-tree: A practically efficient and worst-case optimal R-tree L Arge, MD Berg, H Haverkort, K Yi ACM Transactions on Algorithms (TALG) 4 (1), 1-30	323	2008
The priority R-tree: A practically efficient and worst-case optimal R-tree L Arge, M De Berg, HJ Haverkort, K Yi Proceedings of the 2004 ACM SIGMOD international conference on Management of ...	323 *	2004
The buffer tree: A new technique for optimal I/O-algorithms L Arge Workshop on Algorithms and Data structures, 334-345	310	1995
Scalable sweeping-based spatial join L Arge, O Procopiuc, S Ramaswamy, T Suel, JS Vitter VLDB 98, 570-581	294	1998
The buffer tree: A technique for designing batched external data structures L Arge Algorithmica 37 (1), 1-24	215	2003

refine by coauthor

- Pankaj K. Agarwal (26)
- Jeffrey Scott Vitter (24)
- Ke Yi ⁰⁰⁰¹ (12)
- Laura Toma (11)
- Gerth Stølting Brodal (9)
- Norbert Zeh (8)
- Kasper Green Larsen (7)
- Morten Revsbæk (7)
- Thomas Mølhav (6)
- Jan Vahrenhold (6)
- Octavian Procopiuc (6)
- Peyman Afshani (5)
- Jeff Erickson ⁰⁰⁰¹ (5)
- Herman J. Haverkort (5)
- Jakob Truelsen (4)
- Jeffrey S. Chase (4)
- Rajiv Wickremesinghe (4)
- Constantinos Tsirogiannis (4)
- Michael T. Goodrich (4)
- Andrew Danner (4)

65 more options

refine by venue

- Symposium on Computational Geometry (12)
- ESA (10)
- Algorithmica (8)
- ALENEX (6)
- WADS (5)
- SODA (5)
- SWAT (4)
- ACM J. Exp. Algorithmics (4)
- FOCS (4)
- ISAAC (4)
- PODS (4)
- SIGSPATIAL/GIS (3)
- SSTD (3)
- SIAM J. Comput. (3)
- BigData (2)
- SPAA (2)
- ACM Trans. Algorithms (2)
- ICDT (2)
- Comput. Geom. (2)
- STOC (2)
- GIS (2)

22 more options

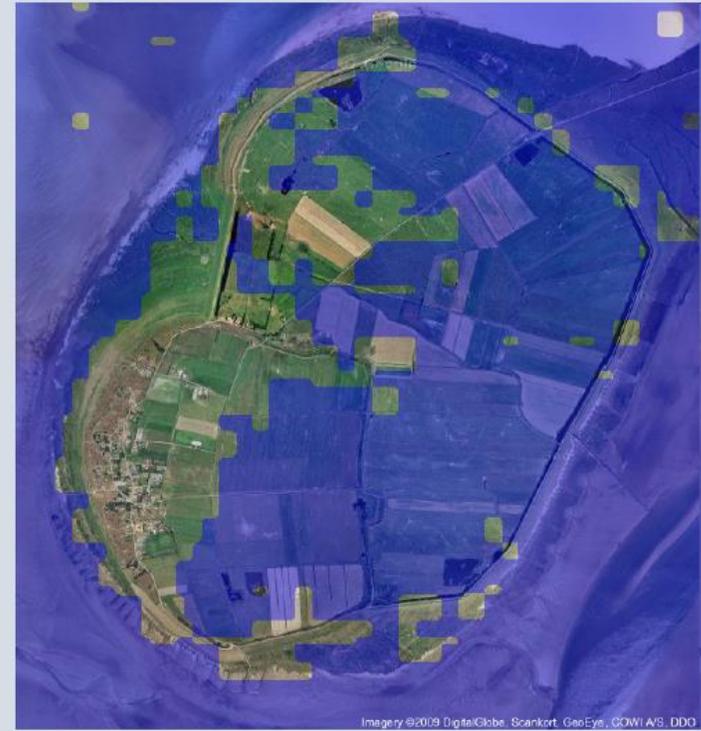
Lars Arge

The Royal Danish Academy of Science and Letters
2016





2 meter resolution model
and 2 meter sea level rise

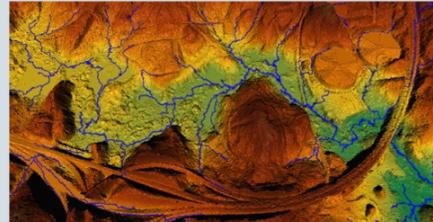


80 meter resolution model
and 2 meter sea level rise

TerraSTREAM Hydrological Tools

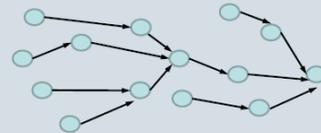
Flow Modeling

Consider a terrain where it rains uniformly across the terrain. The goal of flow modelling is to compute where rain water will accumulate into rivers. The figure to the right shows a computed river network on top of a terrain.

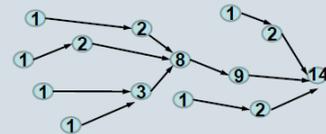


TerraSTREAM can perform flow modelling on arbitrary sized terrain models using provably efficient algorithms. Two of the most important concepts in flow modelling are flow routing and flow accumulation.

Flow routing is the assignment of flow directions to every point of a terrain that models how water flows through it. Flow routing is performed locally at each point by selecting one or more down slope neighbors that receive the incoming flow from upslope neighbors of the point.

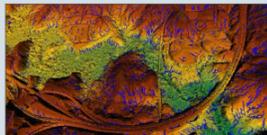


Flow accumulation quantifies how much water flows through each point of the terrain if poured uniformly onto it. Water is routed along the flow paths computed by flow routing.

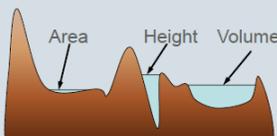


Topological Simplification

Since flow models route water downstream, many small depressions impede water flow.



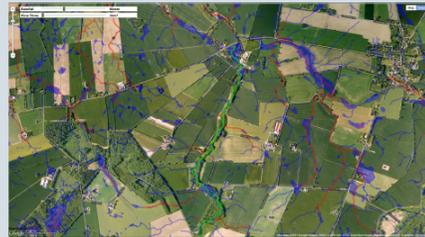
Using topological simplification we can remove depressions using their height, area or volume as a significance measure.



The flow model at the top of the poster has been done on a topologically simplified terrain. Small depressions have been removed so the river network is connected while the big depression in the valley have been kept.

Flash Flood Mapping

Flash flood mapping models how depressions in the terrain fills up and spills into neighboring depressions when it rains.



The output of flash flood mapping shows areas where water will accumulate given a certain amount of rain and which flow paths it follows.

Sea Level Rise Mapping

As a consequence of climate change the sea level is expected to rise significantly around the globe. The TerraSTREAM sea level rise mapping module can map the impact of a rising sea level with provable efficiency on arbitrarily large terrain models.



Shows the parts of Denmark that are flooded when the sea level rises 3.5 meters.

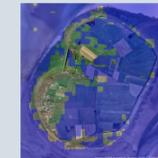
TerraSTREAM has been used to compute sea level rise maps for the Danish government on a terrain model of Denmark consisting of 18 billion points.

Massive Terrain Data is Essential

- Decreasing the size of the data by lowering resolution, will often remove important features such as dikes from the terrain model. The figure to the right shows how the low resolution model fails to capture the dikes around Mandø.
- Computing on smaller parts of the data will not work since flooding can occur from the boundary of the data



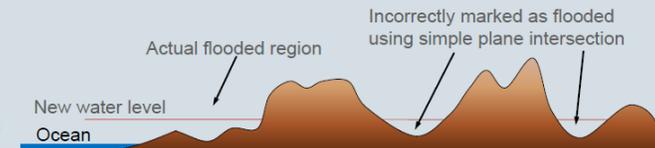
2 meter resolution model and 2 meter sea level rise



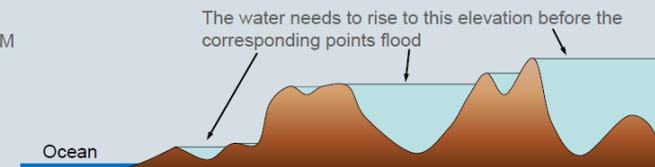
80 meter resolution model and 2 meter sea level rise

Computing the Sea Level Rise Map

The standard way to compute sea level rise maps on bigger terrains is to simply mark all points below a certain height. This does *not* work in general as shown in the figure on the right.



TerraSTREAM computes a new DEM where the elevation of each point in the terrain is the height the water needs to rise to in order for that particular point to flood in the real terrain.



SCALGO

Startup established in late 2009 to commercialize I/O-efficient technology for efficient processing of massive terrain data.



Morten Revsbæk (PhD 2014)
Co-founder
Chief Executive Office



Lars Arge
Co-founder
Chairman of the Board



Thomas Mølhave (PhD 2009)
Co-founder
Chief Technology Officer



Pakaj K. Agarwal
Co-founder

2009 – 25th SoCG (with sausages and beer)



MASSIVE Workshop 2010 with SoCG in Snowbird



2013 Summer School on Data Structures



2014 Summer School ...



2015 MADALGO Retreat





Lars Arge with Michael T. Goodrich at an NFL playoff game during SODA 2018 in New Orleans



Lars Arge

* 8.10.1967 – † 23.12.2020

All honour to his name

Æret være hans minde