

Scalable Computation of Topological Abstractions for Scalar Data

Michael Will¹, Jonas Lukasczyk¹, Julien Tierny², Attila Gyulassy³, Gunther Weber⁴, Hamish Carr⁵, and Christoph Garth¹

¹ RPTU University Kaiserslautern-Landau

² CNRS and Sorbonne Université

³ University of Utah School of Computing and the Scientific Computing and Imaging Institute

⁴ Lawrence Berkeley National Laboratory

⁵ University of Leeds



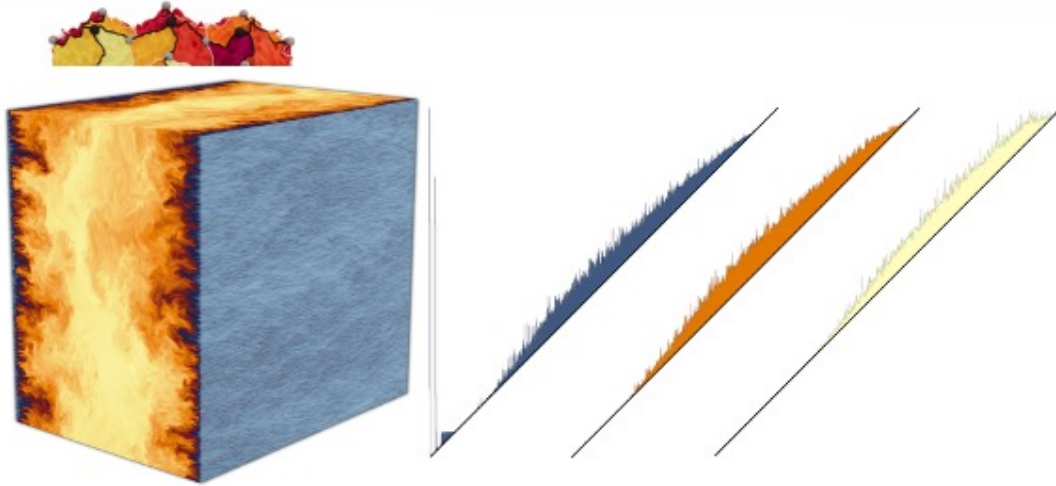
Motivation



- Topological Data Analysis uses concepts from Computational Topology and has become a popular tool for capturing the inherent structure and features of interest of scalar field data
- there are several topological abstractions of interest widely used for many applications
- growing dataset sizes makes computing them on single machines infeasible
- MS complex is a fundamental topological abstraction which segments the domain into areas of similar gradient flow


Motivation

LE GUILLOU E., FORTIN P., TIERNY J.:
Distributed Discrete Morse Sandwich: Efficient Computation of
Persistence Diagrams for Massive Scalar Data, IEEE Transactions on
Parallel & Distributed Systems




- Persistence diagram on a subset of Turbulent Channel Flow with 6 billion vertices, computed with over 4TB of RAM

Motivation



- segmented merge tree on the entropy field of a Richtmyer-Meshkov instability simulation
- Fluid dynamics simulations can scale to hundreds of thousands of processors


Motivation



- Extracting filaments and voids for cosmological density distributions
- The resulting topological abstraction may be simple, but the underlying data may be arbitrarily complex, Nyx cosmological hydrodynamics solver with terabytes of data

Motivation

- no general overview over different parallelization and distribution techniques exist
- Previous surveys (Heine et al. 2016, Yan et al. 2021) have focused on definitions and applications of TDA, but not the algorithmic scalability
- TDA captures *global* structure, but parallelism requires *local* processing



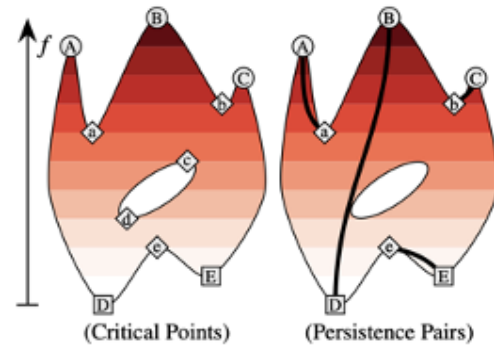
- no general overview over different parallelization and distribution techniques exist
- TDA captures *global* structure, but parallelism requires *local* processing

A very brief introduction to Computational Topology

...at least the parts relevant for our STAR

A very brief introduction to Computational Topology

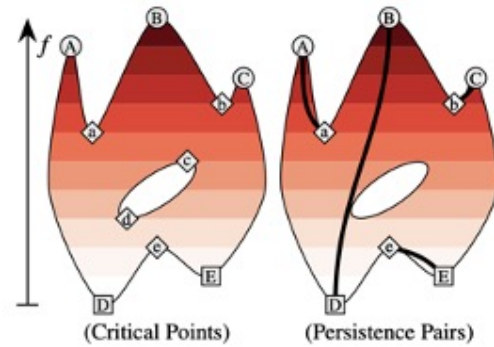
- A scalar field assigns a value to every point in a domain, e.g. elevation in a mountain range
- **Critical points** are where the topology changes: minima (valleys), maxima (peaks), saddles (passes)
- As we sweep a threshold up through the data, connected components appear, merge, and disappear



- Data: piecewise linear scalar fields with data at vertices of d-dimensional connected simplicial complex. Values at edges and faces are interpolated

A very brief introduction to Computational Topology

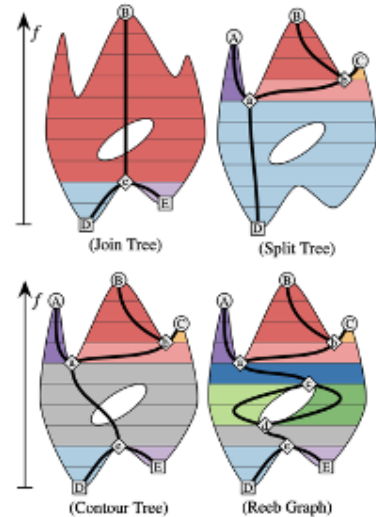
- Each feature is born at a critical point and dies when it merges with an older feature
- **Persistence** is how long a feature lives, measuring its importance
- Persistence Pairs and diagrams show the birth-death pairs



Elder rule

A very brief introduction to Computational Topology


- **Merge Trees:** track when connected components of sub/superlevel sets appear and merge
- **Contour trees:** track connectivity of level set contours on simply-connected domains, combination of upper and lower merge tree
- **Reeb graphs:** generalize contour trees to domains with holes and handles



- Sublevelset: all vertices which have a lower function value than an isovalue
- Reeb graphs: non simply connected


A very brief introduction to Computational Topology

- **Stable and unstable manifold:** partition the domain by gradient flow behavior, combine into the **Morse-Smale Complex**




Overview

- survey paper about scalable topological data analysis
- 158 papers surveyed
- scalable computation models:
 - serial
 - shared memory parallel
 - distributed memory parallel
- relevant abstractions:
 - Set-based (Persistence Pairs / Diagrams)




Overview

- survey paper about scalable topological data analysis
- 158 papers surveyed
- scalable computation models:
 - serial
 - shared memory parallel
 - distributed memory parallel
- relevant abstractions:
 - Set-based (Persistence Pairs / Diagrams)
 - Graph-based (Merge Tree, Contour Tree, Reeb Graph)



Overview

- survey paper about scalable topological data analysis
- 158 papers surveyed
- scalable computation models:
 - serial
 - shared memory parallel
 - distributed memory parallel
- relevant abstractions:
 - Set-based (Persistence Pairs / Diagrams)
 - Graph-based (Merge Tree, Contour Tree, Reeb Graph)
 - Complex-based (Morse-Smale Complexes)



Common building blocks

- Many algorithms share the same low-level components, a "toolkit" for parallel TDA:

- **Union-Find**: maintains connected components; parallelized via path compression and pointer jumping; used in some algorithms for all abstractions



Common building blocks

- Many algorithms share the same low-level components, a "toolkit" for parallel TDA:


- **Union-Find**: maintains connected components; parallelized via path compression and pointer jumping; used in some algorithms for all abstractions



Common building blocks

- Many algorithms share the same low-level components, a "toolkit" for parallel TDA:


- **Union-Find**: maintains connected components; parallelized via path compression and pointer jumping; used in some algorithms for all abstractions



Common building blocks


- Many algorithms share the same low-level components, a "toolkit" for parallel TDA:

- **Union-Find**: maintains connected components; parallelized via path compression and pointer jumping; used in some algorithms for all abstractions



Common building blocks


- Many algorithms share the same low-level components, a "toolkit" for parallel TDA:
- **Union-Find**: maintains connected components; parallelized via path compression and pointer jumping; used in some algorithms for all abstractions
- **Data-parallel primitives**: scan (prefix sum), sort, MapReduce; form the backbone of data-parallel algorithms like Parallel Peak Pruning




Input	1	2	3	4	5
Prefix Sum	1	3	6	10	15

Common building blocks

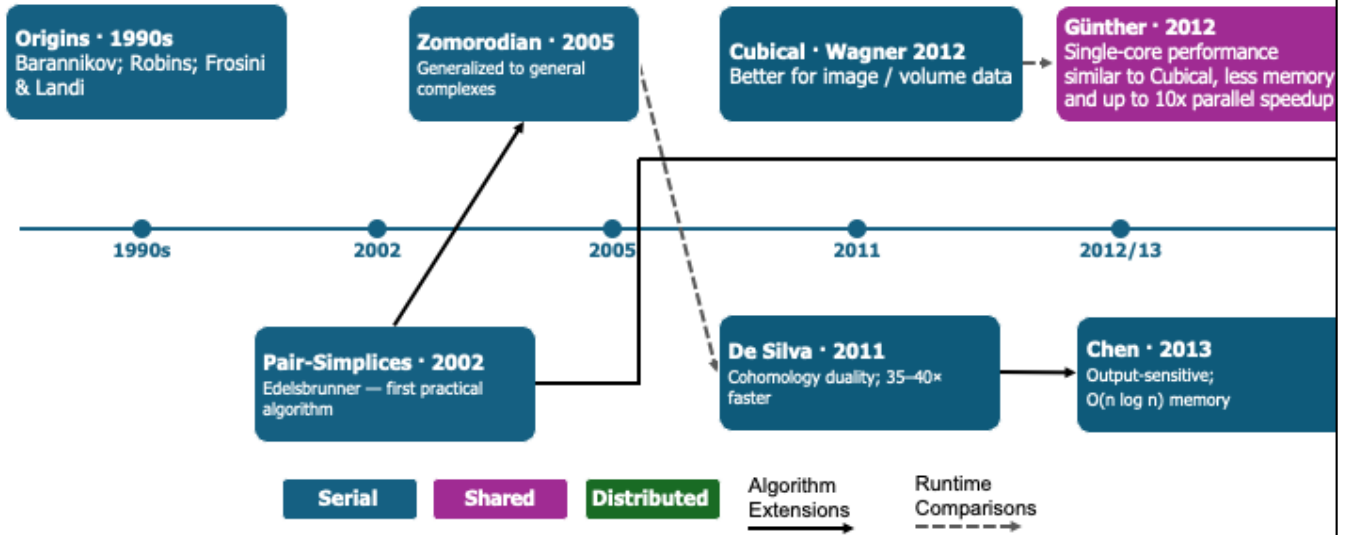
- Many algorithms share the same low-level components, a "toolkit" for parallel TDA:
- **Union-Find**: maintains connected components; parallelized via path compression and pointer jumping; used in some algorithms for all abstractions
- **Data-parallel primitives**: scan (prefix sum), sort, MapReduce; form the backbone of data-parallel algorithms like Parallel Peak Pruning
- **Discrete gradient**: embarrassingly parallel to compute; each vertex processed independently; crucial for Discrete Morse Theory-based approaches



Input	1	2	3	4	5
Prefix Sum	1	3	6	10	15

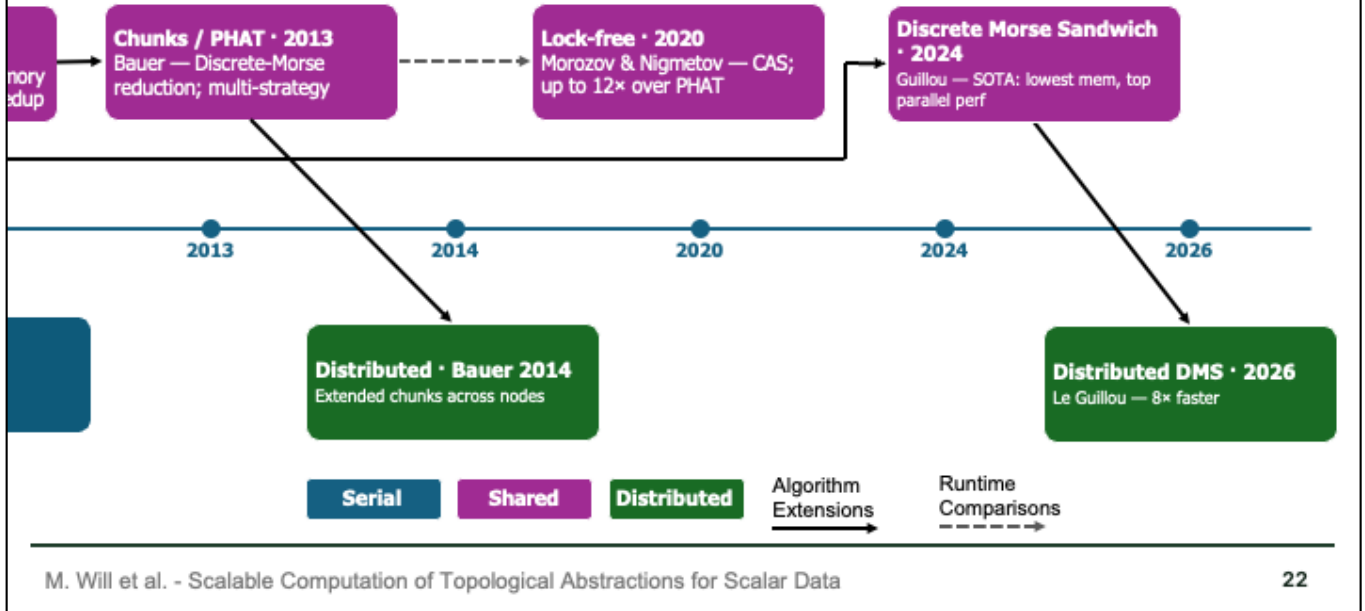


Persistence: Serial Foundations



- Persistence independently introduced by barannikov, robins and frosini&landi which partition critical values into birth death pairs by upper triangular matrix reduction
- Pair simplices captures and quantifies topological features across different filtration levels in a simplicial complex. also formally introduced the persistence diagram
- While the worst-case complexity of this algorithm is cubic, in practice it is significantly faster
- Zomorodian generalized it for higher dimensional data
- One branch: De Silva et al have presented a method exploiting dualities between homology and cohomology to accelerate persistence diagram computation
- This led to the first output sensitive algorithm by Chen et al. which solves previous reduction on rank computation on submatrices. Do not improve worst case complexity, but strong results, especially in memory
- Other branch: cubical complexes instead of triangular complexes

Persistence: Parallel



- Morozov: compare and swap
- DMS: combining Pair-Simplices with discrete morse theory and including aggressive specializations for 3D. strong memory and time improvements

Merge Trees: Serial Foundations

- Originally used for radar landscape simulation (Boydell & Ruston 1963)
- Intuition: imagine flooding a landscape: islands appear at maxima, merge at saddles
- Carr et al. (2003): standard algorithm using Kruskal + Union-Find in $O(n \log n + t\alpha(t))$ time
- Merge trees are a lightweight alternative to the full contour tree, especially useful for in-situ analysis

- **Kruskals minimum spanning tree**
- **n vertices, t simplices, α is the inverse Ackermann function which grows very very slowly**

Merge Trees: Shared Memory

- **FTM-Tree** (Gueunet et al. 2017): task-parallel Kruskal; propagations start at all maxima in parallel, synchronize at saddles; up to 8× faster than serial; memory-heavy for large sweep fronts

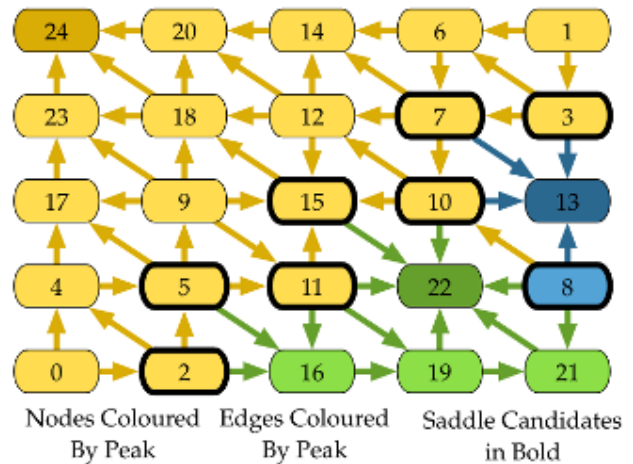


- major drawback of this approach is maintenance of the entire sweep front using Fibonacci heaps, which require a lot of memory for large sweep fronts in large datasets
- Another unintended consequence of starting the tasks at the maxima is that the tasks are not evenly distributed, which can lead to load imbalances and lower parallel efficiency when increasing the size of the dataset but not the amount of the actual work done e.g. when resampling an initial dataset to a higher resolution for scaling experiments. While the merge tree may not increase in size, FTM-Tree will still require substantially more memory and time to compute the merge tree on the higher resolution dataset.

Merge Trees: Shared Memory

CARR H. A., WEBER G. H., SEWELL C. M., AHRENS J. P.:
Parallel peak pruning for scalable SMP contour tree computation.
LDAV 2016


- **FTM-Tree** (Gueunet et al. 2017): task-parallel Kruskal; propagations start at all maxima in parallel, synchronize at saddles; up to 8× faster than serial; memory-heavy for large sweep fronts
- **Parallel Peak Pruning** (Carr et al. 2016): data-parallel using sort and scan primitives; iteratively prunes leaf branches; up to 20× faster than FTM-Tree in 3D, Parallel Peak Pruning has various extensions for shared and distributed memory



- Instead of vector based approach, PPP uses data parallel primitives, identifying saddle candidates in the data by checking which monotone paths reach multiple peaks

Merge Trees: Shared Memory

LUKASCZYK J., WILL M., WETZELS F., WEBER G. H., GARTH C.:
ExTreeM: Scalable Augmented Merge Tree Computation via
Extremum Graphs. IEEE TVCG 2023



- **ExTreeM** (Lukasczyk et al. 2023): first derives the extremum graph (orders of magnitude smaller than the domain), then runs merge tree computation on that; up to 2 orders of magnitude faster than FTM-Tree and up to 11x faster than PPP

Merge Trees: Distributed Memory

- **Morozov & Weber** (2013): local-global representation where each processor stores a sparse version of the global tree; up to 14× faster than parallel contour tree computation
- **Werner & Garth** (2021): task-parallel distributed algorithm growing regions around critical points; up to 5× faster than Morozov & Weber with better scaling
- **Distributed merge forest** (Huang et al. 2021): localized trees connected by bridge sets; up to 14.6× faster workflows for large-scale queries

-Key use case: in-situ analysis of combustion simulations at volume size 2025×1600×400 on 30,000 cores

Contour Trees

- **Serial:**

- Carr et al. (2003): established the standard $O(n \log n)$ algorithm
- Raichel & Comandur (2022): first algorithm without a global sort


- **Shared memory:** Contour Forests (Gueunet 2016): partition domain by value, compute local trees, stitch; up to 4× faster than serial

- **Distributed:**

- Li et al. (2024): first fully distributed augmented contour tree; up to 2 orders of magnitude faster than serial;
- Li et al. (2025) adds pre-simplification for even larger datasets at 2.58× speedup and 6.67× lower memory

- Idea of Carr: merge the join tree and split tree (two merge trees) in a postprocessing step


Research Gaps



Research Gaps



Research Gaps



- **Reeb graphs:** burst of serial papers 2006–2015 bringing serial computation to maturity; only 1 paper after 2015; zero dedicated distributed algorithms
- **Morse-Smale complex:** heavy 2006–2015 activity; only 4 papers after 2015; distributed work can only handle simplified (not full) complexes
- **No unified benchmark:** surveyed works span decades and different hardware, a controlled reimplementa-tion-based benchmark remains an open and valuable contribution

Open Problems and Research Gaps

- Distributed Reeb graph algorithms: contour tree algorithms partially fill the gap on simply-connected domains, but the general case has no solution
- Distributed Morse-Smale complex: current approaches (Gyulassy et al. 2012; Will et al. 2024) require local simplification or only compute segmentations, not the full complex
- Potential path forward: extend local-global representations (used successfully for merge/contour trees) to Reeb graphs and MS complexes
- ML-based TDA methods are evolving rapidly but currently lack exactness guarantees, possibly a future survey topic once the area matures

Conclusion

- TDA is powerful, but only useful at scale if we can compute it at scale
- Three-tier taxonomy: serial → shared memory → distributed memory, across five major abstractions
- Common building blocks (Union-Find, data-parallel primitives, discrete gradient) underpin many scalable algorithms
- Strong recent progress in persistence and merge/contour trees; significant gaps remain for Reeb graphs and MS complexes in distributed settings
- A unified benchmark across reimplemented methods on standardized hardware remains the most needed next contribution

- we hope that our work can guide future work by highlighting the state of the art and the gaps needing to be filled

Further information in the paper

- Applications

Further information in the paper

- Applications
- Performance comparisons

Abstraction	Authors	Paper	Performance
Counter Trees	Pavesio et al.	Parallel Computation of the Topology of Level Sets [PCMS04]	parallel speedup-walks almost linearly
	Arhava et al.	A parallel and memory efficient algorithm for enumerating the counter tree [ARH15]	better performance and better parallel speedup and using only one fifth of memory compared to [PMS02]
	Chen et al.	Counter Trees: Fast multi-threaded augmented counter trees [CST14]	up to 4x faster than [CST15], up to 10x faster than [CST14], up to 1.5x faster than [CST15], up to 1.5x faster than [CST14]
	Sharma et al.	On-Demand Augmentation of Counter Trees [OS20]	up to 4x faster than [CST15], up to 1.5x faster than [CST14], up to 1.5x faster than [CST15], up to 1.5x faster than [CST14]
Rank Graph	Chen et al.	Fast-Seed Augmented Rank-Graphs with Dynamic W-Trees [CST17-18]	up to 1x faster than [CST15] using our formal, 9.76x faster using 32 threads
	Wu et al.	An Efficient Data Reduced Parallel Rank Graph Algorithm [WR20]	up to 1.5x faster than [CST15], up to 1.5x faster than [CST15]
MI Counting	Chen et al.	Memory Efficient Computation of Persistence Homology for 3D Images Using Dynamic Morse Theory [EMDF17]	1.5x with a critical point, scales only linearly with number of vertices, especially faster on noisy data than [PMS14]
	Strothmann et al.	Efficient computation of 3D Morse-Smale complexes and persistent homology using dynamic Morse theory [SRW15]	up to 20x speedup [SRW15], up to 20x faster than [PMS14]
	Strothmann et al.	Parallel Computation of 3D Morse-Smale Complexes [SMN13]	near linear parallel scaling
	Cygan et al.	Parallel Computation of 3D Morse-Smale Complexes [SMN13]	order of magnitude faster than [SRW15] and [EMDF17], and less memory used
	Cygan et al.	Parallel Computation of Morse-Smale Complexes with Improved Accuracy [CST14]	10.7x faster than [CST15], slower than [SRW15] and [EMDF17], but more accurate
	Strothmann et al.	GPU Parallel Computation of Morse-Smale Complexes [SRW15]	10x faster than [CST15], up to 10x faster than [CST15], up to 10x faster than [CST15]
	Maack et al.	Parallel Computation of Piecewise Linear Morse-Smale Segmentation [MST15]	up to 1.5x faster than [CST15], up to 1.5x faster than [CST15]

Abstraction	Authors	Paper	Performance
Persistence	Bauer et al.	Distributed Computation of Persistence Homology [BAK15]	comparable to [BAK15], but well scaling distributed memory usage
	Lewis et al.	Parallel Computation of Persistence Homology using the Steiner Complex [LM15]	well performance scaling, but high memory usage where distributed data not help
Merge Trees	Le Gallier et al.	Distributed Dynamic Morse Smale Complex: Efficient Computation of Persistence Diagrams for Massive Scalar Data [LGP19]	comparable in time to [LGP19], but well scaling distributed memory usage, on average 1x faster than [LGP19]
	Munier et al.	Distributed merge trees [MNF15]	up to 10x faster than [PMS15]
	Lindig et al.	In-Situ Feature Extraction of Large Scale Continuum Simulations Using Segmented Merge Trees [LJC14]	comparable to [MNF15], but better scaling to higher core counts
	Kiciorczyk et al.	Toward Localized Topological Data Structures: Querying the Forest for the Tree [KGF20]	one order of magnitude faster than [PST19] and [PST14], without scaling distributed memory usage
	Werner et al.	Overlaid Task Parallel Augmented Merge Tree Construction [WET11]	up to 1x faster than [MNF15] and [LGP14], with better scaling
	Huang et al.	Distributed merge forest: a new fast and scalable approach for topological analysis of mesh [HAF15]	not computing the whole tree, but faster queries when counter trees are large, up to 1.6x faster work than [MNF15]
Counter Trees	Barker et al.	3D-Counter Queries and Gradient Descent with Guaranteed Delivery in Sensor Networks [BDF16]	no comparison
	Munier et al.	Distributed Counter Trees [MNF14]	not computing the whole tree, but faster queries, good distributed scaling
	Nath et al.	Minimally parallel algorithms for computing TIN DEMs and counter trees for large datasets [NCS14]	distributed MPC model versus using [PMS15]
	Blanc et al.	W-Structures in Counter Trees [BC12]	abstract theory on W-structures in parallel counter trees
	Carlier et al.	Distributed Hierarchical Counter Trees [CKN11]	significantly faster than [PMS15]
	Li et al.	Distributed Augmentation, Hypergraphs, and Branch Decomposition of Counter Trees for Scientific Datasets [LCK12]	up to 2 orders of magnitude faster than initial [PMS15]
		Extremely Scalable Distributed Computation of Counter Trees via Pre-Simplification [LCK12]	better scaling and up to 1.6x faster than [LCK12] with up to 6.7x faster counter usage

Further information in the paper

- Applications
- Performance comparisons
- Existing software packages

Software Package	Abstraction	Parallelization Type	Publications	Link	Language
YTK	Persistent Homology Merge Trees Contour Trees Rect Graphs MS Complex Merge Trees	Shared Memory Distributed Memory (Persistent Homology)	[TUL16,MBP21,LOW24]	https://topology-tool-kit.github.io/	C++ Python (Bindings)
Nikarn	Merge Trees	Shared Memory Distributed Memory	[MSU*16,MBP21,MBP23]	https://www.nikarn.org/	C++
PersistentCycles	Persistent Homology	Shared Memory	[W21]	https://github.com/tact1c10f/persistentcycles	YTK Plugin C++ Python (Bindings)
CaBi	Persistent Homology	Shared Memory	[MBW11]	https://caBi.sta.fu-berlin.de/	C++
Empire	Persistent Homology	Serial	[Mw17]	https://www.mpi-nwz.de/software/empire/	C++
Phi	Persistent Homology	Shared Memory	[BKRW17]	https://github.com/3f3n0x/phi	Python (Bindings) C++
EDHA	Persistent Homology	Shared Memory Distributed Memory	[BKR17]	https://github.com/3f3n0x/edha	C++
Irregular	Persistent Homology	Serial	[NY18]	https://github.com/irregular-topology/irregular	Java Matlab
Tracer	Persistent Homology	Serial	[BKS17]	https://github.com/3f3n0x/tracer	Java
Quinn	Persistent Homology	Serial	[SRSW24]	https://github.com/AppleLiu/Quinn	C++
Fornax	Persistent Homology	Serial	[NS13,NS21]	http://pqdla.math.ubc.ca/fornax/persistent	C++
Onco	Persistent Homology	Shared Memory	[DND1]	https://github.com/antipater/onco/	C++ Python (Bindings)
Rather	Merge Trees	Shared Memory Distributed Memory	[PNS18]	https://github.com/rtz/rather	C++
libsort	Merge Trees	Serial	[JRS15,ImpkowskaCS18]	https://github.com/ivab11ant/libsort	C++
CaBi2	Persistent Homology	Shared Memory Distributed Memory	[Wag15,Wag16]	https://github.com/HubWag/caBi2	C++
MSComplex	MS Complex	Shared Memory	[MBF18]	https://github.com/ser111/MSComplex	C++
Microplex	MS Complex	Shared Memory	[SLL12,SHY13]	https://github.com/ser111/microplex	C++

