David Tench — High-Performance Algorithms Engineer

CONTACT INFORMATION

Email: dtench [at] pm [dot] me
Website: www.davidtench.com
Github: https://github.com/tenchd
LinkedIn: https://www.linkedin.com/in/david-tench/

SKILLS

I design and build high-performance systems for graphs and databases that use randomized algorithms to massively scale up and scale out applications in domains like bioinformatics and databases.

PROFESSIONAL EXPERIENCE

Lawrence Berkeley National Lab, Berkeley, CA

Grace Hopper Postdoctoral Fellow — **June 2023 – June 2026**

Rutgers University, New Brunswick, NJ

CRA Computing Innovation Postdoctoral Fellow — **Aug 2021 – June 2023**

SKILLS

- Design and analysis of massive-scale algorithms and data structures
- Expertise in graph algorithms/systems, out-of-core computation (distributed, memory-hierarchy-optimized), randomized algorithms, parallelism
- Writing high-performance Rust & C++ code. Ex: https://github.com/tenchd/spec_spars
- Profiling and diagnosing performance of high-performance implementations
- Leading a high-performance algorithms coding team (17 students)
- Technical writing: 12 peer-reviewed papers, 4 grant applications
- Technical presentation: 18 conference or workshop talks

SELECTED PROJECTS

Linear Sketching for Large-Scale Graph Processing

Team Lead, Algorithm Designer

- Noticed 1) that existing graph processing systems fail to scale to massive, dense, dynamic graphs and 2) algorithms in theory literature (called **graph sketches**) that might help, but had never been implemented due to high computational and data movement costs.
- Designed new randomized graph sketching algorithms that simplify and reduce the computational cost of the SOTA, and optimize for out-of-core (distributed, SSD) computation.
- Recruited and led 8-student team to implement sketching algorithms through GraphZeppelin, a high-performance multi-core graph processing system. C++14, OpenMP, MPI.
- 1st author paper appeared at SIGMOD, world's top-ranked database CS research conference.
- 1st author paper, a distributed version (increasing throughput by more than 10x), appeared at ALENEX, a top-ranked algorithmic engineering CS research conference.
- Papers currently in submission to top-ranked conferences expand the above ideas to GPU acceleration, fast dynamic connectivity queries, and hybrid sketching algorithms. See project repos at https://github.com/GraphStreamingProject.

Randomized Graph Algorithms for Compacting Memory Management

Algorithm Designer/Engineer

- Designed real-time randomized matching graph algorithms used to reduce C++ application memory footprint by compacting objects without changing virtual addresses.
- Experimentally verified runtime and solution quality of novel algorithms.
- Co-1st author paper appeared at PLDI, a top-ranked operating systems conference.
- Mesh, our C++ compacting memory manager: https://github.com/plasma-umass/Mesh

EDUCATION

University of Massachusetts Amherst, Amherst, MA

Ph.D., Computer Science and M.S., Computer Science — **Aug 2014 – Aug 2020**

Dissertation title: "Algorithms for Massive, Changing, and Otherwise Inconvenient Graphs"

Lehigh University, Bethlehem, PA

B.S., Mathematics — **Aug 2009 – May 2013**