Jonathan Derryberry

http://www.cs.cmu.edu/~jonderry/

Education

2003 - 2009	Carnegie Mellon University Computer Science Department Ph.D. in Computer Science Advisor: Daniel Sleator Area: Data Structures and Algorithms
1998 - 2003	Massachusetts Institute of Technology M.Eng. in Electrical Engineering and Computer Science (GPA: 5.0/5.0) B.S. in Computer Science and Economics with a minor in Mathematics (GPA: 4.8/5.0)

Graduate Coursework: Theory of Computation, Advanced Complexity Theory, Cryptography, Advanced Cryptography, Algorithms, Randomized Algorithms, Approximation Algorithms, Theory of Performance Modeling, Electronic Marketplaces, Machine Learning, Networks, Type Systems, Advanced AI Concepts, Computer Graphics, Optimizing Compilers, Mathematical Games

Publications

- **Adaptive Binary Search Trees**
  Ph.D. Thesis CMU-CS-09-180, Carnegie Mellon University, 2009.

  Committee: Daniel Sleator (chair), Guy Blelloch, Gary Miller, and Seth Pettie

  `http://www.cs.cmu.edu/~jonderry/thesis.pdf` (current version)
  `http://www.cs.cmu.edu/~jonderry/slides.pdf` (slides)

  Introduced a new binary search tree algorithm, cache-splay, that was the first to achieve the Unified Bound of Iacono since this problem was posed in 2001, and showed that the lower bound framework (see below) applied to a more general model of computation than binary search. Additionally, my thesis summarized my previous work on developing both lower bounds and upper bounds on the cost of sequences of accesses in the binary search tree model.

- **Properties of Multi-Splay Trees**
  Technical Report CMU-CS-09-171, Carnegie Mellon University, 2009.

  Collaborators: Daniel Sleator and Chengwen Chris Wang

  `http://www.cs.cmu.edu/~jonderry/mst_properties.pdf`

  Showed that multi-splay trees (see below) satisfy even richer adaptive properties than those that were shown in the original multi-splay tree paper. In particular, we showed that they satisfy the working set bound, and that they perform as well as a double-ended queue (deque) when used as one.

- **Skip-Splay: Toward Achieving the Unified Bound in the BST Model**
  In *Proceedings of the 11th International Symposium on Algorithms and Data Structures (WADS 2009)*, pages 194–205, Berlin, Heidelberg, 2009. Springer-Verlag.

  Collaborators: Daniel Sleator

  `http://www.cs.cmu.edu/~jonderry/skip-splay.pdf`

  Developed skip-splay, the first binary search tree algorithm to provably achieve good adaptive performance in comparison to the Unified Bound, which essentially states that accesses are cheap whenever they are near a recent access. Skip-splay is extremely simple and almost identical to splaying.

- **Achieving Spatial Adaptivity While Finding Approximate Nearest Neighbors**
  In *Proceedings of the 20th Canadian Conference on Computational Geometry (CCCG 2008)*, pages 163–166, 2008.

  Collaborators: Don Sheehy, Daniel Sleator, and Maverick Woo

  `http://www.cs.cmu.edu/~jonderry/annfinger.pdf`

  Developed the first approximate nearest neighbor data structure that was provably capable of exploiting spatial proximity of queries to achieve faster performance than that which is achieved by running the standard worst-case data structure.

- **Experimental Evaluation of Parametric Max-Flow Algorithms**
  In *Proceedings of the 6th Workshop on Experimental Algorithms (WEA 2007)*, pages 256–269, 2007.

  Collaborators: Maxim Babenko, Andrew Goldberg, Robert Tarjan, and Yunhong Zhou

  `http://www.cs.cmu.edu/~jonderry/maxflow.pdf`

  Compared two algorithms for bipartite parametric maximum flow, one with good worst-case guarantees and one that often exhibited good performance in practice. The emphasis of this work was on finding the characteristics of easy problem instances by generating various sequences of instances that exposed gaps in the comparative performance of the two algorithms.

- **O(log log n)-Competitive Dynamic Binary Search Trees**
  In *Proceedings of the 17th ACM-SIAM Symposium on Discrete Algorithms (SODA 2006)*, pages 374–383, New York, NY, USA, 2006. ACM.

  Collaborators: Chengwen Chris Wang and Daniel Sleator

  `http://www.cs.cmu.edu/~jonderry/multi-splay.pdf`

  Introduced multi-splay trees, the first binary search tree data structure proved to be $O(\log \log n)$-competitive while supporting not just queries, but insertions and deletions as well. We also showed that multi-splay trees had a variety of other adaptive properties.

- **A Lower Bound Framework for Binary Search Trees with Rotations**
  Technical Report CMU-CS-05-187, Carnegie Mellon University, 2005.

  Collaborators: Daniel Sleator and Chengwen Chris Wang

  `http://www.cs.cmu.edu/~jonderry/lower-bound.pdf`

  Generalized previous instance-specific lower bounds for the binary search tree model by developing a lower bound framework that consisted of an optimization problem whose solution was always a lower bound on the cost of the optimal binary search tree for the particular sequence of queries.

- **Combinatorial Auctions with Structured Item Graphs**
  In *Proceedings of the 19th National Conference on Artificial Intelligence (AAAI 2004)*, pages 212–218. AAAI Press / The MIT Press, 2004.

  Collaborators: Vincent Conitzer and Tuomas Sandholm

  `http://www.cs.cmu.edu/~jonderry/AAAI104ConitzerV2.pdf`

  Introduced the following results relating to the tractability of combinatorial auctions: showed how to efficiently solve a special case of the winner determination problem, showed how to efficiently detect whether a set of bids contained a particular structure that made them easy to solve, and showed that another restriction of combinatorial auctions was NP-complete.

- **Creating a Web Page Recommendation System for Haystack**
  M.Eng. Thesis. 2003. Massachusetts Institute of Technology.

  Advisor: David Karger

  `http://www.cs.cmu.edu/~jonderry/derryberry_masters.pdf`

Built a website recommending tool for the Haystack project. The recommender was capable of using the text of articles, the hierarchical structure of the URLs, and the layout position of the link to learn which new links a user was likely to be interested in. This functionality, as well as a training data gathering mechanism was fully integrated into the Haystack information management system.

- **Exploring Simon Kirby's Parameter Space**
  Undergraduate Thesis. 2002. Massachusetts Institute of Technology.

  Advisor: Patrick Winston

  `http://www.cs.cmu.edu/~jonderry/AUP_writeup.pdf`

  Wrote an application that facilitated the study of the evolution of grammatical languages among simulated agents as per Simon Kirby's theory of the evolution of language without natural selection. Performed an array of simulations with varied parameter values to expand upon Kirby's findings. The trend in the results showed the parameters used by Kirby were at a sharp transition point between disparate evolutionary patterns.

Additional Projects

- (Spring 2004) Built software to track disease outbreaks using Markov random fields as part of an original group project for the class Advanced AI Concepts.

  Collaborators: Matt Streeter

  `http://www.cs.cmu.edu/~matts/cs15780/finalproj/`

- (Fall 2002) Built software for converting individual two-dimensional images into three-dimensional models by using the variation in the image's brightness as part of an original group project for the class Computer Graphics.

  Collaborators: Jeremy Cannon and Vitaly Kulikov

  `http://groups.csail.mit.edu/graphics/classes/6.837/F02/projects/reports/team13.pdf`

Work Experience

2006	**Summer Intern, HP Labs** Worked on the bipartite parametric max-flow problem. Designed experiments to test the performance of competing algorithms, one with good theoretical guarantees and one with good performance on many instance-types that appear in practice. Also worked on developing algorithms for easier special cases.
2004	**Summer Intern, IBM Research** Worked on variants of the resource-constrained project scheduling problem for applications to the steel industry.
2001	**Summer Intern, LifeHarbor Investments, Inc.** Developed the client-side portion of a three-tier Java application to be used as a tool that facilitates the maintenance of managed accounts.
2000	**Summer Intern, Shym Technology** Member of the QA team for a computer security application.
1999	**Summer Intern, Vector Marketing Corp.** Sales representative for Vector, sellers of Cutco cutlery.

Teaching Experience

Fall 2007	**Volunteer Teaching Assistant for Advanced Data Structures** `http://www.link.cs.cmu.edu/15859-f07/` Gave some lectures and participated in the development of the course.
Spring 2006	**Teaching Assistant for Graduate Algorithms** `http://www.cs.cmu.edu/afs/cs/academic/class/15750-s06/www/index.html` Gave a few lectures and review sessions, devised some homework and exam problems, and graded assignments and exams.
Fall 2005	**Volunteer Teaching Assistant for Advanced Data Structures** `http://www.cs.cmu.edu/afs/cs/academic/class/15859-f05/www/index.html` Gave some lectures and participated in the development of the course.
Fall 2004	**Teaching Assistant for Undergraduate Algorithms** `http://www.cs.cmu.edu/afs/cs/academic/class/15451-f04/www/` Taught weekly recitations in which lecture material was reviewed and some new material was introduced. Graded written and oral assignments and exams.

Skills

Java, C++, ML, Scheme, Matlab, Mathematica, LaTeX, LEDA, Java3D, Weka, Unix, Windows.

Interests

Puzzle solving. TopCoder. Online poker. Nutrition and fitness. MIT Crew member during freshman year. Alumnus of Kappa Sigma. Recreational basketball and running.

Personal Details

Name	Jonathan Derryberry
Email	jonderry@cs.cmu.edu
Telephone	(412) 708-1509
Web Page	`http://www.cs.cmu.edu/~jonderry/`
Address	5000 Forbes Avenue Computer Science Department Pittsburgh, PA 15213
Nationality	American

References

- Daniel Sleator, Professor of Computer Science, Carnegie Mellon University: sleator@cs.cmu.edu
- Tuomas Sandholm, Professor of Computer Science, Carnegie Mellon University: sandholm@cs.cmu.edu