



# From the Lab to the Data Center: My Experiences Commercializing Science

---

Anindya Datta  
Walking Stick Solutions



# Agenda

---

- Background
- Journey into entrepreneurship
- Current status



# Background

---

- After completing my Ph.D. in 1994 (U. Maryland), I spent several years in academia
  - First at the Univ. Arizona (1993-1998)
  - Then at Georgia Tech (1998-2003)
- I've always had a desire to commercialize my research, to prove to myself that my research has utility in the real world. To this end, I've started three companies:
  - Nanosystems (1994)
  - Chutney Technologies (2000)
  - Walking Stick (2007)
- I'd like to give you a sense of the journey this has been



# Nanosystems: History

---

- I was working in active, real-time, data-driven systems
  - Specific interest in developing systems that could compute derived values very efficiently
  - Great potential for applicability in the financial sector
- Built a working prototype with 2 grad students
- Started looking for ways to commercialize



# Nanosystems: Technology

---



# Nanosystems: Related Journal Publications

---

- A. Datta and S.H. Son, A Study of Concurrency Control in Real-Time Active Database Systems, *IEEE Transactions on Knowledge & Data Engineering*, 14(3), 2002.
- A. Datta and I. Viguier, Handling Sensor Data in Rapidly Changing Environments to Support Soft Real-Time Requirements, *INFORMS Journal on Computing*, 12(2), 2000
- A. Datta, S.H. Son, V. Kumar, Is A Bird in Hand Worth More than Two In the Bush? Limitations of Priority Cognizance in Conflict Resolution for Firm Real-Time Database Systems, *IEEE Transactions on Computers*, 49(5), pp 482-502, May 2000
- A. Datta, S. Mukherjee and I. Viguier, PAPER: Prefetching and Priority Cognizant Buffer Replacement Policies in Real-Time Active Database Systems, *Journal of Systems and Software*, 42, 1998, 227-246
- A. Datta and I. Viguier, A Novel Approach for Smaller, Faster Data Warehouses, *DATABASE*, 29(4), pp 33-44, 1998
- A. Datta and I. Viguier, Providing Real-Time Response, State Recency and Temporal Consistency in Databases for Rapidly Changing Environments, *Information Systems*, 22(4), 1997, pp 171-198
- A. Datta and A. Ghosh, Analyzing Concurrency Control and Scheduling Policies in a Real-Time Database Systems Using the Framework of Discrete Event Dynamical Systems, Automata and Formal Language Theories, *Foundations of Computing and Decision Sciences*, 22(4), 1997, pp 221-250
- A. Datta, S. Mukherjee, P. Konana, I. Viguier and A. Bajaj, Multiclass Transaction Scheduling and Overload Management in Firm Real-Time Database Systems, *Information Systems*, 21(1), January 1996, pp. 29-54




# Nanosystems: Related Conference Publications

---

- A. Datta and I. Viguier, Buffer Management in Active, Real-Time, Database Systems, Proc. Workshop on Active, Real-Time Database Systems, Como, Italy, September 1997
- A. Datta, S. Chakravarthy, S. Thomas and I. Viguier, An Architecture and Two New Research Problems in ARCS Databases, DART (Databases: Active and Real-Time) Workshop, November 1996
- A. Datta, A. Bajaj, A High Performance, Low Cost, Priority Cognizant Concurrency Control Algorithm for Firm Real-Time Database Systems, Proceedings of the First INFORMS Conference on Information Systems and Technology, May 1996
- A. Datta, Databases for Active Rapidly Changing data Systems (ARCS): Augmenting Real-Time Databases with Temporal and Active Characteristics, Proceedings of the First Workshop on Real-Time Database: Issues and Applications, Newport Beach, March 1996.

# Nanosystems: Outcome and Lessons Learned

---





# Chutney: History

---

- I was working on novel indexing techniques, and (with colleagues), had developed a unique method of storing data and indexes
  - Trouble publishing in the big conferences (SIGMOD/VLDB) – the paper always got one “clear accept – loved it” review and one “hated it” review
- I thought the idea had some commercial merit, so with 3 grad students, started to develop software based on it



# Chutney: Technology

---

- We started out developing data warehousing software
  - Quickly found that we could not hope to compete with the big vendors (Oracle and many others)
  - So we could not get funding for the company
  - We were considering developing a high-performance recommendation system to make use of the technology
- At around the same time, dynamic content generation technologies for web sites came into wide use
  - Web sites became mighty slow
  - We decided to try to build software to improve web site performance through caching
- The technology evolved over time into generic object storage and virtualization software



# Chutney: Related Journal Publications

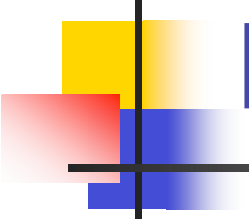
---

- "Optimization in Object Caching", by K. Dutta, S. Soni, S. Narasimhan, A. DATTA, *INFORMS Journal on Computing*, 18(2): pp 243-254
- K. Dutta, A. Datta, D. VanderMeer, P. Keskinocak, Optimal Caching Strategies for E-Commerce Sites, *IEEE Transactions in Systems, Man and Cybernetics*, Forthcoming
- D. VanderMeer, A. Datta, K. Dutta, H. Thomas, K. Ramamritham, Proxy-Based Acceleration of Dynamically Generated Content on the World Wide Web, *ACM Transactions on Database Systems*, 29(2): 403-443, June 2004.
- D. VanderMeer, A. Datta, K. Dutta, K. Ramamritham, S.B. Navathe. Mobile User Recovery in the Context of Internet Transactions. *IEEE Transactions on Mobile Computing*, 2(2), 2003.
- A. Datta, K. Dutta, H. Thomas, D. VanderMeer, World Wide Wait: A Study of Internet Scalability and Cache-Based Approaches to Alleviate It, *Management Science*, 49(10): 1425-1444, October 2003.
- A. Datta, D. VanderMeer, K. Ramamritham. Parallel Star Join + DataIndexes: Efficient Query Processing in Data Warehouses and OLAP. *IEEE Transactions on Knowledge and Data Engineering*, 14(6), 2002.
- A. Datta, K. Dutta, H. Thomas, D. VanderMeer, K. Ramamritham, D. Fishman. A Comparative Study of Alternative Middle Tier Caching Solutions to Support Dynamic Web Content Acceleration. *IEEE Internet Computing*, 6(5), 2002
- A. Datta, K. Dutta, D. VanderMeer, K. Ramamritham, S. Navathe, An Architecture to Support Scalable Online Personalization on the Web, *VLDB Journal*, 10(1), 2001
- H. Thomas, A. Datta, A Conceptual Model and Algebra for On-Line Analytical Processing in Decision Support Databases, *Information Systems Research*, 12(1), 2001

# Chutney: Related Conference Publications



- D. VanderMeer, H. Thomas, K. Dutta, A. Datta, K. Ramamritham, ReDAL: Request distribution for the Application Layer. Proceedings of the International Conference on Distributed Computing Systems (ICDCS), Columbus, OH, June 2005
- A. Datta, K. Dutta, H. Thomas, D. VanderMeer, Suresha, K. Ramamritham, Proxy-Based Acceleration of Dynamically Generated Content on the World Wide Web: An Approach and Implementation, Proc. ACM SIGMOD, June, 2002
- A. Datta, K. Dutta, H. Thomas, D. VanderMeer, K. Ramamritham, D. Fishman, A Comparative Study of Alternative Middle Tier Caching Solutions to Support Dynamic Web Content Acceleration, Proc. 27th VLDB Conference, September, 2001
- A. Datta, K. Dutta, K. Ramamritham, H. Thomas, D. VanderMeer, Dynamic Content Acceleration: A Caching Solution to Enable Scalable Dynamic Web Page Generation, Proc. ACM SIGMOD, May, 2001
- A. Datta, K. Ramamritham and H. Thomas, Curio: A Novel Solution for Efficient Storage and Indexing in Data Warehouses, Proc. 25th VLDB Conference, Edinburgh, Scotland, September, 1999
- A. Datta and I. Viguier, A Novel Approach for Smaller, Faster Data Warehouses, Proc. Workshop on Information technologies and Systems 1997 (WITS'97), Atlanta, December 1997



# Chutney: Outcome and Lessons Learned

---

- The technology currently serves as a part of the foundation for Cisco's AON (application-oriented networking) product



# Walking Stick: History

---

- I am currently interested in how applications access data – patterns, problems, interactions
- With the Oracle-PeopleSoft merger and Larry Ellison's accompanying promises to stop supporting PeopleSoft as soon as possible, I thought there might be a business opportunity in helping PeopleSoft customers move to other enterprise platforms
  - Oracle's plans to stop supporting the platform, however, have gradually lost their teeth – PeopleSoft has several years of promised support at this time
- In further talking with enterprise software users (business and technical folks), I started to think about where my expertise and interests could be applied to help better manage and run enterprise installations




# Walking Stick: Technology

---

- A major problem in enterprise installations is the gap in visibility between the business service (the part that provides the business value) and the actual execution of the coded logic on the hardware and software platform
  - Business services are long-running and distributed
    - May run over the course of multiple days, with long stretches of inactivity
    - Touch multiple machines and applications on those machines
  - Hundreds of business services may be running at a given time
- When problems happen, it's difficult to determine which low-level process is the culprit, and even harder to determine the bottom-line business impact of all the other business services that are delayed or unexpectedly halted
- We are building monitoring software to support better visibility between the high-level business services and the low-level resources and data they access

# Walking Stick: Current Status and Lessons Learned So Far





# Final Thoughts

---

- There's a thread running through this description, implying that moving from the lab to the market is hard, and not for the faint of heart
- Yes, many things have gone wrong on my journey, often spectacularly so
- You may wonder, then, why am I still doing this?