

DISTINGUISHED LECTURE SERIES
EECS COLLOQUIUM
Fall 2009



Task-specific Search

Michael Stonebraker

M.I.T. and Goby, Inc

Abstract

Historically, search engines have focused on searching the entire web for objects of interest using a single “one-size-fits-all” keyword query system supported by word-oriented indexing. In a task-specific search application such as real estate values (Zillow.com), airline tickets (Kayak, et. al..) and comparison shopping, such general purpose information retrieval (IR) techniques can be drastically improved. One can leverage any or all of the following to generate a vastly improved search experience.

- ▶ a task-specific user-interface for queries
- ▶ a task-specific user interface for displaying results
- ▶ a task-specific data model for semantically organizing data
- ▶ a curation mechanism for deciding which sites have potentially valuable data
- ▶ a mechanism for searching such web sites and converting returned data to this data model (semantic data integration)
- ▶ a mechanism for searching for data that is behind HTML forms (the so-called deep web)
- ▶ a task-specific ranking system to order the presentation of returned data
- ▶ a task-specific mechanism for caching popular (static) returned data

This talk describes a task-specific search engine, Morpheus, built at M.I.T. and explains how it solves the above issues. Then, we describe the changes that Goby made to Morpheus to produce a commercial product. We conclude with a demo of the Goby product.

Biography

Dr. Stonebraker has been a pioneer of data base research and technology for more than a quarter of a century. He was the main architect of the INGRES relational DBMS, and the object-relational DBMS, POSTGRES. These prototypes were developed at the University of California at Berkeley where Stonebraker was a Professor of Computer Science for twenty five years. More recently at M.I.T. he was a co-architect of the Aurora/Borealis stream processing engine, the C-Store column-oriented DBMS, the H-Store transaction engine, and the Morpheus search engine. Each has been commercialized by a venture-capital backed startups. Presently he serves as Chief Technology Officer of Vertica, which is commercializing C-Store, VoltDB which is commercializing H-Store, and Goby which is commercializing Morpheus.

For the full biography, please go to:

<http://www.eecs.berkeley.edu/Colloquium/Archives/09-10/Fall2009/stonebraker.shtml>

Wednesday
October 28
4:00 - 5:00 pm

306 Soda Hall
Hewlett-Packard
Auditorium

SPONSORED BY A GENEROUS DONATION FROM **HEWLETT-PACKARD**