ARCHIVAL RESEARCH · ADITYA MOHAN

Decentralized Data Discovery at HP Labs: Kundali

Research era: 2002 (research appointment) · Originally published on Blogger: 2012-10-31 · Archive recovered October 2026

> **2026 editorial context (not part of the original post).** Concise original summary of Kundali, distributed data summaries, query routing and caching, with linked 2002 technical report and 2003 proceedings.
>
> 2026 context: illustrates work on decentralized state and retrieval well before Bitcoin. This was not a Bitcoin-like cryptographic ledger.

Original Blogger address: http://adiinsf.blogspot.com/p/peer-to-peer-information-retrieval-hp.html (live availability not established). Old report links below are carried forward but not guaranteed to work.

Original 2012 archive text

HP LABS | Palo Alto, CA USA (6/2002–9/2002) Distributed Search and Caching

(Research was done while I was a computer science graduate student at Brown University)

Information retrieval from decentralized, self-organizing network. My research on speculative routing and update propagation focused on the second face of the problem " Given the exact reference of an object of interest, how can you efficiently retrieve the correct object from a fully decentralized, self-organizing network". The answer was Kundali data structures that maximized the probability to find the information, minimize peer access latencies and balance the workload among many peers. Application for this research was for Collaborative Peer-to-Peer infrastructure for Personalized Services. The key challenge in peer-to-peer collaborative environment is how to efficiently locate and retrieve the correct data. Techniques for efficient searching in peer-to-peer networks have been proposed; however, these handle location and routing as a single problem and impose a structure in the network by mapping the data to particular nodes. In the project, we proposed propagation and routing algorithms for a fully decentralized, self-organizing network. Our goal was to maximize the probability of finding the data, minimize peer access latencies and balance the workload among many peers. Central to our approach was the Kundali data structure that represents the set of data maintained by the peers and drives the smart routing of the search requests (queries). Kundali, for each peer, maintained a Bloom Filter based set of synopsis

of the data expected to be present at each routing direction. Requests that cannot be answered locally, were propagated only to those immediately connect peers whose synopsis depict the closest match. The algorithm was implemented in the context of a fully decentralized Internet caching service. Related publication “Speculative Routing and Update Propagation: A Kundali Centric Approach”. This work was one of the first few research in the world on how to use peer-to-peer systems for providing personalized services.

Related Publications

- Aditya Mohan, Vana Kalogeraki, “Speculative Routing and Update Propagation: A Kundali Centric Approach” in IEEE 2003 International Conference on Communications, May 11-15, 2003, Anchorage, AK, USA (**www**) [box.net folder]
- Aditya Mohan, Vana Kalogeraki, “Speculative Routing and Update Propagation: A Kundali Centric Approach”, HP Labs 2002 Technical Report, HPL-2002-237, August 22, 2002 (**www**) [box.net folder]
- Vana Kalogeraki, Aditya Mohan, "A Caching Architecture for Peer-to-Peer Systems", ACM/IFIP/USENIX International Middleware Conference (Middleware 2003), Rio de Janeiro, Brazil, June 16-20, 2003. [www] [box.net folder]