



Research Activities of the PDSG Group

Computing Systems Laboratory

<http://www.cslab.ece.ntua.gr/pdsg>

National Technical University of Athens

Greece

General info

People:

- ❑ Nectarios Koziris (Associate Professor, NTUA)
- ❑ 4 post-doc researchers
- ❑ more than 15 graduate students



Research areas

- High performance computing
 - Challenging Applications
 - Lack of inherent (easy) parallelism (e.g. Dijkstra)
 - Memory bandwidth saturation (e.g. SPMxV, Floyd-Warshall)
 - Memory latency (graph algorithms)
 - Frequent synchronization
 - Application impact on architecture
 - Architecture impact on applications
 - Different architectures (PC clusters, CMPs, GPGPUs, Cell B/E)



Research areas

- Computer architecture
 - Memory Hierarchy (eg. Cache Partitioning)
 - SMT (eg. Thread Synchronization, Resource Sharing)
 - Exploration and evaluation of state-of-the-art architecture designs



Research areas

- **Distributed Systems**

- P2P computing+Grid computing+Cloud computing
- Focus on Data Management
- Solutions for storing, indexing and transferring data in structured and unstructured overlays
- Adaptive resource location and query processing
- Load balancing



P2P Facts

- The Internet has three valuable assets
 - Information
 - Computing resources
 - Bandwidth
- All are vastly under-utilized due to the traditional client-server model!!
- Bandwidth consumption up to 60% (!!) of the total Internet traffic
- The size of the networks and the complexity/ requirements increase
 - P2P for email, web-caching, IM, scientific collaborations, telephony, etc



Why P2P?

- Cost sharing/reduction
 - Commodity hardware
- Improved scalability
- Reliability
- Resource aggregation
- Increased autonomy
- Self-organization



Ongoing Work

- HiPPIS [HPDC'09, P2P'08, WIDM'08, CoopIS'08]
- Brown Dwarf [Subm. to EDBT'10]
- GridTorrent [FGCS'09, DOA'09]



HiPPIS

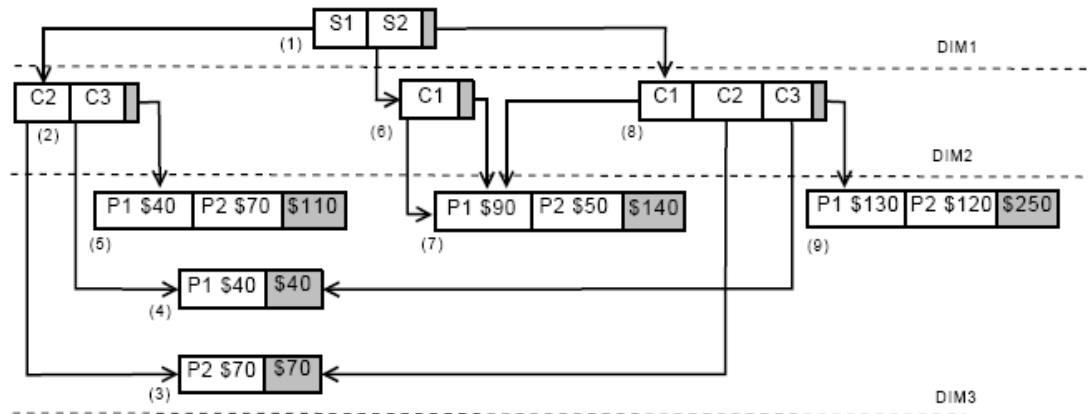
- Complete system for storing and querying hierarchical data in DHTs
- Initial insertion at a pivot level (combination)
- Point and group-by queries answered through
 - DHT lookups when concerning pivot level combination
 - Soft-state indices, created upon query misses
- Adaptive re-indexing according to current query trend
 - Per peer local statistics for incoming queries
- Online Updates



Brown Dwarf

- Based on Dwarf structure
 - computes, stores, indexes and updates fully (or not) materialized cubes
 - Eliminate prefix and suffix redundancies

DIM1	DIM2	DIM3	Measure
S1	C2	P2	\$70
S1	C3	P1	\$40
S2	C1	P1	\$90
S2	C1	P2	\$50

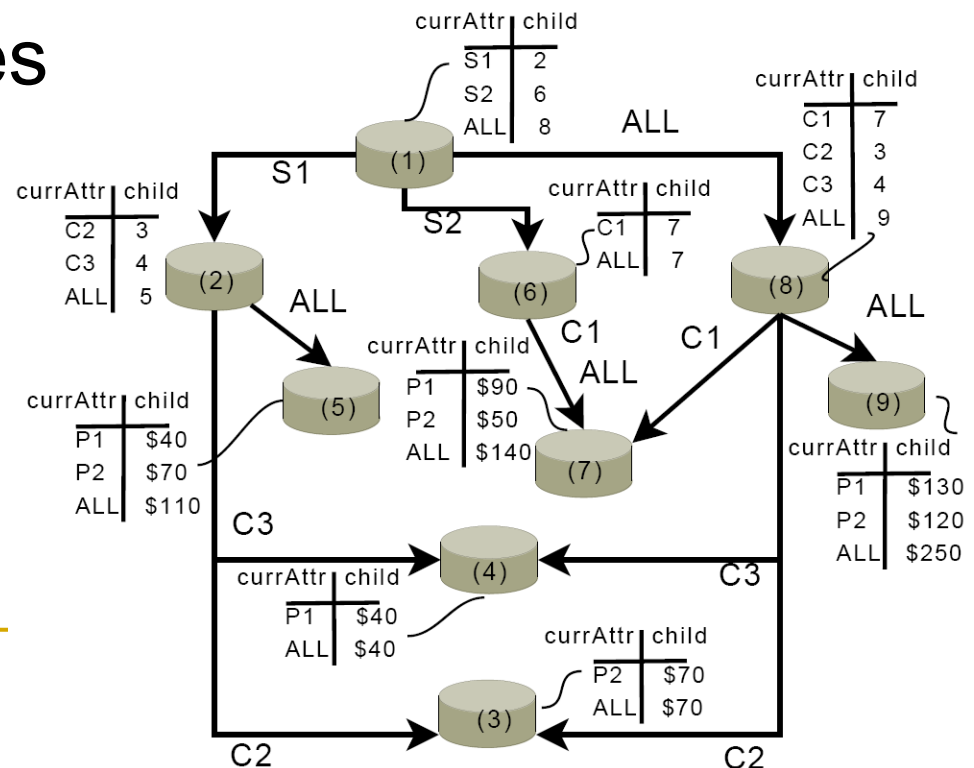


- Why distribute it?
 - Store even larger amounts of data!
 - Dwarf reduces but may also blow-up data (sparse cubes)
 - Handle many more requests
 - Update and query the system online



Brown Dwarf

- Dwarf nodes mapped to overlay nodes
- One-pass over the fact table insertion
- Queries → Overlay path of d hops
- Incremental Updates
- Mirroring (k)
 - Load
 - Churn



GridTorrent

- Implementation of the popular BitTorrent protocol
- interface with Data Grid components and protocols (e.g. GridFTP, RLS)
- Allows clients to download files from multiple sources while uploading them at the same time
- Combine the best out of the two protocols
 - Fragment selection provides optimized data transfer service.
 - Striped version of GridFTP protocol



Looking for ...

- Exchange of research ideas
- Networking with other universities and industry partners
- Collaboration
- Spawn our ideas to real-life projects



Thank you!

- <http://www.cslab.ece.ntua.gr/pdsg/>
- <http://www.cslab.ece.ntua.gr/~doka>

