NTNU
Norway

DASCOSA: Database Support for Computational Science Applications

Kjetil Nørvåg
Norwegian University of Science and Technology
Trondheim, Norway

Outline

- Background/context: Databases & Grids
- Requirements for Grid-wide database systems
- Peer-to-peer technology
- DASCOSA framework
- Project info: Organization and research approach
- Summary: Expected results

First: Why use databases instead of files?

Advantages of databases:

- Support durability
- Manage concurrent accesses to data
- Efficient data access
- Maintenance of metadata together with associated data
- Application←→data independence

- Lesson learnt in other application areas years ago… Developers of computational science applications still haven't…?

The Grid

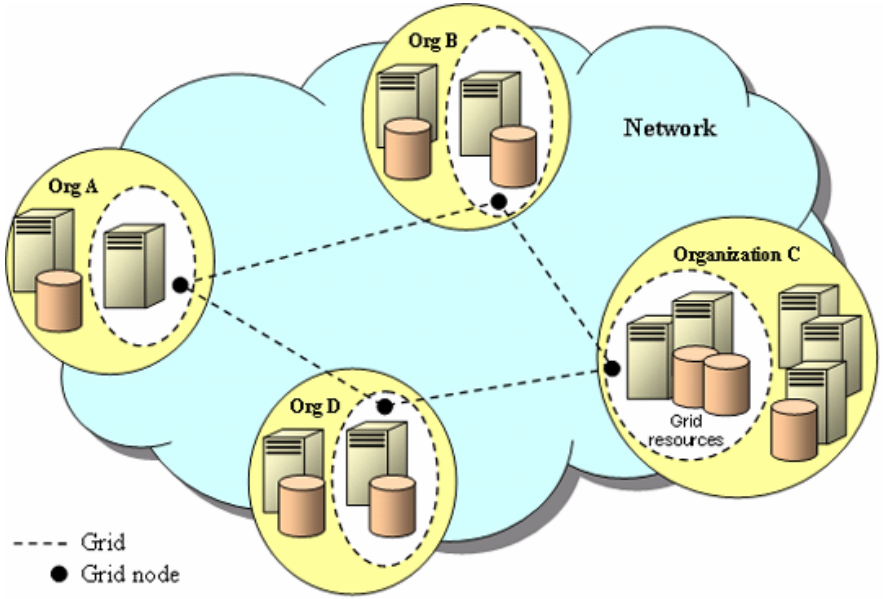

- The original Grid concept: Location-transparent computation
- Nowadays: lot's of different definitions ☺

  Three point check-list for a grid (Ian Foster, ANL):

  1) Coordinates resources that are not subject to centralized control

  2) Uses standard, open, general-purpose protocols and interfaces

  3) Delivers non-trivial qualities of service
- Currently: Mostly stable organizations and hardware
- However: Will change for future Grid/eScience applications!

→ Want self-organization!

The Data Grid

- Grid computing gaining maturity
- Management of *data* in Grid systems less mature
- Data Grid: Mostly file oriented
- Some support for metadata management

- Emerging standard proposals/frameworks for more advanced services:
  - Examples: OGSA-DAI and OGSA-DQP
- Typical for proposals: Service-based, no cooperation between sites

- Our goal: *Database Grid* with:
  - Location-transparent storage: don't have to care about where data is stored and where queries are processed!
  - Cooperation on processing while retaining autonomy
- I.e., *Grid-wide DBMS!*

Requirements for Grid-wide DBMS

- Durability
- High availability
- Security
- Scalability
- Performance
- Hardware/OS independent
- Metadata management & resource discovery

Application-wise relaxation of requirements possible

Tool in satisfying requirements: Peer-to-peer (P2P) techniques

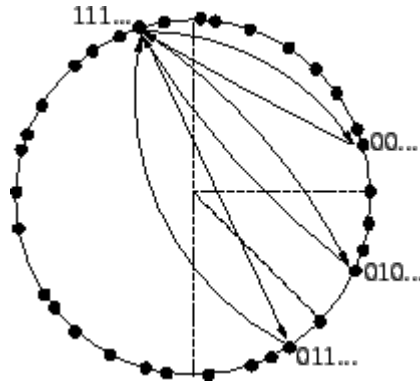

- What is P2P?
  - Peers act as equals, merging the roles of clients and server
  - There is no central server managing the network
  - One of main differences with pure Grid: Creation and maintenance of overlay network with structure independent of underlying internet
- Advantages:
  - Scalability
  - Robustness
  - Less administrative cost
- P2P already used in a number of applications, e.g.:
  - File sharing (Kazaa, Bittorrent...)
  - Telecommunication: Skype

DASCOSA architecture: P2P-based DBMS

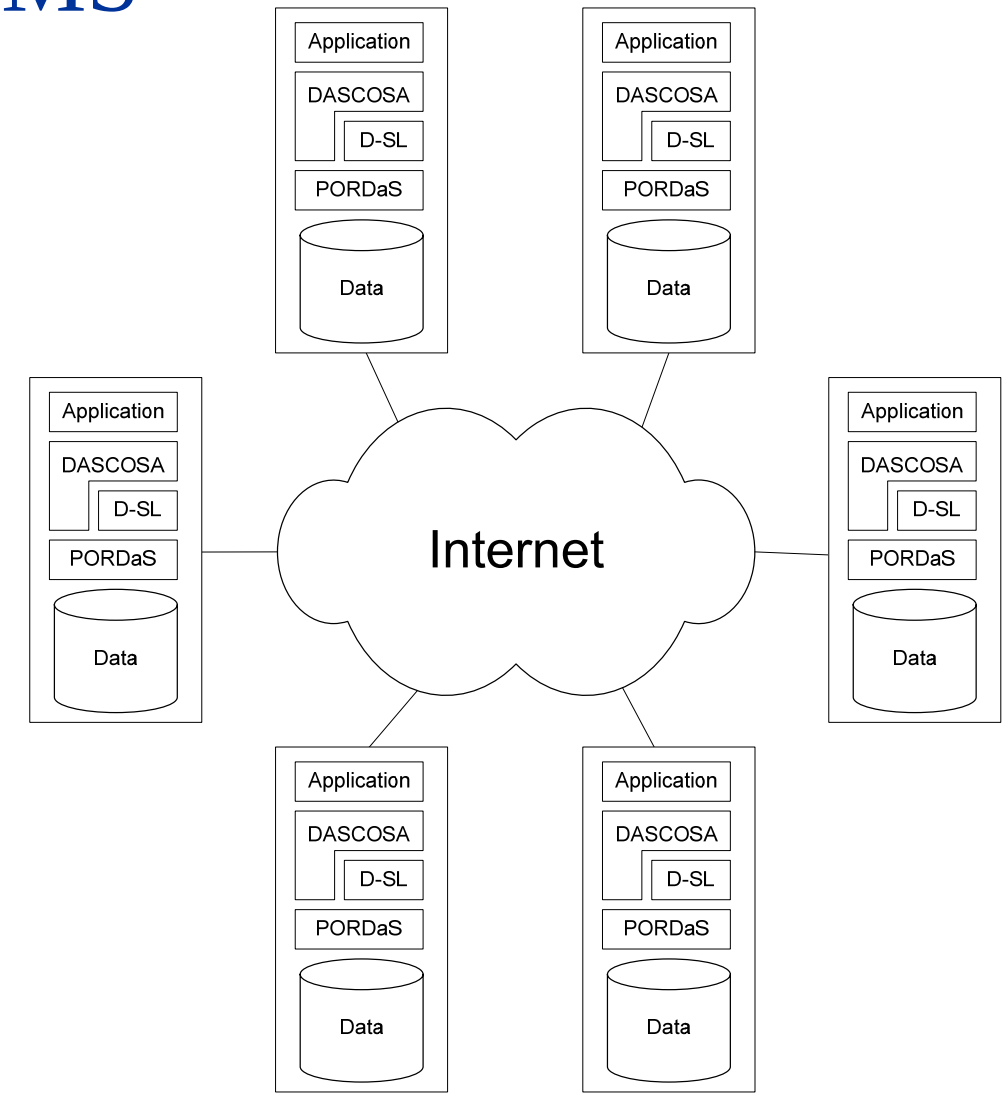

Basic middleware: PORDaS (*P*2P *O*bject-*R*elational *Da*tabase *S*ystem)

- Based on P2P technology
- Object-relational model
- Data created & stored locally but might be made globally available for querying
- Support for both replication and caching of data
- Interesting challenges:
  - Query planning
  - When to cache/replicate, when to use, which to use...
  - How to efficiently index data in such systems
  - ”Participation economics”: Get as much back as you pay

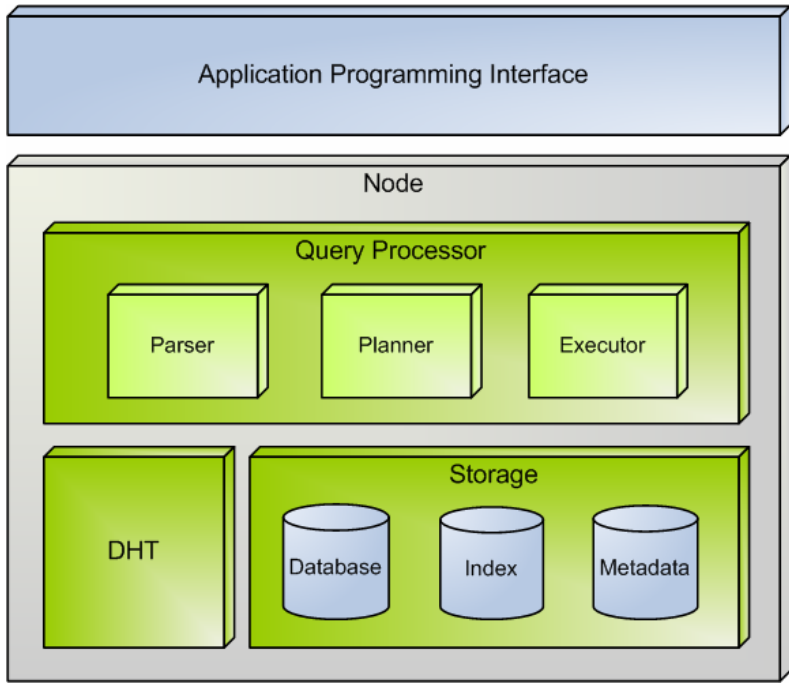

DASCOSA semantic layer: Ontology-based metadata management

- Will play a major role in supporting queries over related data sources using heterogenous schema descriptions
- Database tables can be described by:
  - Table name and attributes names (traditional)
  - Keywords/additional describing text (preferably from standardized taxonomy)
- Data discovery tool needed ← Challenge!
- Will build on Semantic Web/Grid technologies

Approach/work packages

- Pre-project: Proof-of-concept prototype (finished 2006)
- Development of the overall architecture & individual parts of architecture
- Development of appropriate mechanisms for making it possible to satisfy desired requirements, including scalability, durability, and performance
  - Includes algorithms for both storage and data access/querying
- System implementation to show feasibility of the approach
- Performance evaluation of system, possibly extended with analytical modeling to prove characteristics as scalability

Project organization



- Duration: 2007-2010
- Project members:
  - Prof. Kjetil Nørvåg (Database Group), project leader
  - Prof. Jon Atle Gulla (Information Systems Group)
  - Dr. Jon Olav Hauglid, postdoc, RCN-funded
  - N.N., PhD student, RCN-funded
  - (Norvald Ryeng, PhD student, NTNU-funded)
  - Master students
- International cooperation: AUEB, Athens

Research Environment: Database Group at NTNU

- 7 faculty members, wide expertice in database systems, distributed systems, operation systems, and performance evaluation
- History:
  - Design and implementation of database systems since early 70s
  - Several database machines built in 80s
  - ”World records” ☺ in execution of join & sort operators ca. 1990
  - Spinoff companies include:
    - TechRa (later Kvatro/MaxWare)
    - ClustRa (now part of Sun Microsystems)
- Current active cooperation:
  - Companies: Sun, Fast, Telenor, Yahoo, DNV, ...
  - A number of universities both in Norway and abroad

Expected results

- An approach for distributed and decentralized data management for Grid applications that supports:
  - Traditional database features including durable storage and efficient querying
  - Seamless access to the heterogeneous sources in the data Grid
- PhD candidate and a a number of Master candidates
- Papers in refereed journals and high quality international conferences
- Improving network with international research groups working in database-grid related areas
- And hopefully: *Help in making computational engineering society believe in databases!* ☺