GORDA

Scope

Today's businesses face many challenges and trade-offs when making their core data consistently available to answer increasingly demanding performance and dependability requirements. Consistent distributed data management is a complex problem on today's computing and networking technologies, whether it is a local cluster of machines or a worldwide distributed setting. Current databases systems are at the core of the vast majority of existing information systems. Although anchored on the proven transactional paradigm and supported by more than three decades of research, they are still, in practice, very difficult to replicate and manage in a consistent and scalable manner.

Database replication is therefore a key technology for the long-term competitiveness of today's businesses. Replication technology is faced with new requirements. Firstly, enterprise-wide availability should take into account legacy information systems. Secondly, wide-area support is required not only by large businesses but also by small and medium enterprises, which inevitably raises security issues. Finally, wider applicability of database technology calls for highly scalable Database Management Systems (DBMS).

The resulting heterogeneity, geographical dispersion, and scale have a profound impact on the applicability of replication techniques. The wider applicability of DBMS also calls for an increased concern with costs.

The Gorda project fostered database replication as a means to address the challenges of trust, integration, performance, and cost in current database systems underlying the information society. It standardised architecture and interfaces, and sparked their usage with a comprehensive set of components that are ready to be deployed.

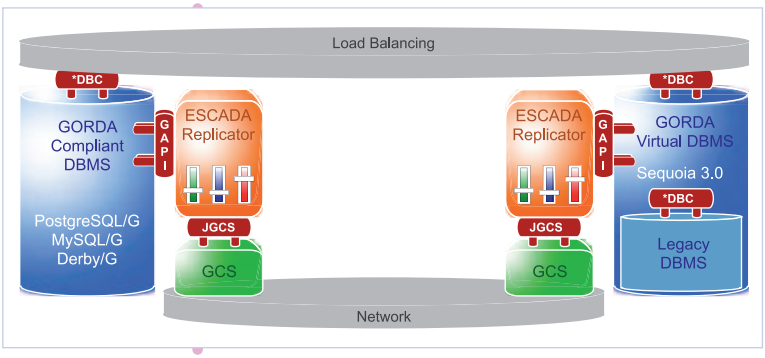

Advances

The Gorda promoted the interoperability of DBMSs and replication protocols. It defined a generic architecture and a set of standard interfaces, resulting in a significant advantage for several market players. Database vendors are now given the means to address concrete requirements of specific markets. End-users are given flexible choices to pick from by not having products tied to a single vendor or proprietary interfaces.

Furthermore, Gorda provides reference implementations for the proposed architecture and interfaces. Many achievements resulting from Gorda activities combined with previous research results have been implemented and prove the adequacy of the project's concepts and technology. The completeness of the Gorda architecture and interfaces was clearly demonstrated. This is a major breakthrough as previous research and industry efforts towards database replication were never tuned concerning product and prototype integration.

Positioning in global context

Gorda has provided fully functional prototypes as proof of concept of the project's main goals. The resulting technology has been mapped into three major databases: MySQL, PostgreSQL and Apache Derby. Additionally, it has also been materialised as an industry product (Sequoia v3.0), commercialised by one of the project partners. Currently, the Gorda replication middleware can be considered the largest and most diversified repository of database replication protocols for existing open source DBMSs.

Contribution to standardization and interoperability issues

Efforts to establish the proposed architecture and APIs as a de facto industry standard were continuously made. The project team has released reference implementations for PostgreSQL, Sun Microsystems Apache Derby and MySQL products, which have triggered discussion, favoured standardisation and resulted in maturation and industry adoption of the proposed interfaces and architecture.

Target users / sectors in business and society

The project's results are directly available to database vendors and software developers. The former should find clear advantage in implementing a replication API such as the one Gorda specifies. By studying the several reference implementations on Open Source DBMSs, such a task should be simple and rewarding, if not straightforward.

Given the availability of replication-aware DBMSs implementing the Gorda API, software developers can leverage many existing replication protocols. Some of them were made publicly available by Gorda. Software developers can focus on tuning and specialising the protocols to specific environments as scenarios. A major advantage is the independence of these protocols with respect to the underlying DBMS.

Finally, system integrators and service providers have now a comprehensive set of alternatives of DBMS and replication protocols to choose from.

Overall benefits for business and society

Database replication is an enabling technology to ensure availability, performance, and geographic distribution of databases. Gorda allows users to construct better database applications. However, Gorda has a more fundamental role. By enabling and improving replication, Gorda also enables scale-out designs, which offer a different economic model from other more capital-intensive approaches to database systems. Scale-out designs work by spreading copies of data and load across many database hosts. This simple mechanism enables users to create highly available and very highly performing systems at a fraction of the cost of other approaches. As requirements increase, users can incrementally add more hardware, thereby scaling the system efficiently.

Scale-out is thus a design with important economic properties that particularly benefit small as well as growing companies. The low initial cost makes such systems accessible to a wide variety of businesses. The fact that capabilities can be scaled incrementally is ideal for growing businesses and fits in well with the utility computing model, where system capacity is scaled up or down based on current needs. It is no accident that countless web businesses use scale-out designs to get started and to provide efficient growth as usage increases. Gorda's focus on replication is thus a root cause for such benefits.

Examples of use

As a simple example of the benefits of Gorda usage, consider a new start-up that is building a social web site. This is a concrete (and common) example of the type of application represented by standard benchmarks, such as TPC-W, an important use case for the Gorda project overall.

Using a scale-out design, the initial version of the site would consist of a single master database with a couple of slave databases to help handle read traffic and to be available as masters should the primary database fail. Gorda provides multiple options for construction of such clusters, but for now we wish to focus simply on the system cost. At current prices such a system can be put together very economically with 3-4 hosts (say € 4000 up-front costs) and a part-time database administrator (say € 15000 yearly operation costs). This is well within the reach of almost all small businesses and has the advantage that as load and availability requirements increase they can be met by adding more hosts. Operational costs are also reasonable as such a system is cheap to host and can be operated by a part-time database administrator who shares other tasks.

By contrast, a traditional approach to designing such a system would involve use of Oracle RAC with a shared disk. Such a system would have up-front costs in the range of $100,000 *minimum* for Oracle licenses plus a shared disk and extra network infrastructure. The additional complexity of such a system would imply more time required to manage the system by an Oracle database administrator, who would also command a higher salary.

For many businesses, the price differences would be so large that the social website would never be built and the business itself would not exist. This same reasoning applies to many other systems both large and small that benefit users and create livelihoods for the persons who operate them.

Achievements

- Gorda Architecture and Programming Interface Specification.
- jGCS, Generic interface for group communication supporting Appia, JGroups and Spread.
- PostgreSQL/G, PostgreSQL Gorda API reference implementation.
- Apache Derby/G, Apache Derby Gorda API reference implementation.
- MySQL/G, MySQL API reference implementation.
- Sequoia/G, Continuent Sequoia Gorda API reference implementation.
- Appia toolkit, Group communication toolkit aimed at efficiently supporting database replication.
- ESCADA replication server, Flexible replication server supporting eager multi-master protocols.
- JadePlus, Autonomic middleware for managing clusters of replicated databases.
- OSGi bundle that allows the OSGi framework to be managed through JMX.
- Bristleconne, Testing tools and benchmarks for scale-out.

title
Open replication of databases

contract number
004758

type of project
Specific Targeted Research Project

contact point
Rui Oliveira
UNIVERSIDADE DO MINHO, PT
e-mail: rco@di.uminho.pt

project website and partner list
http://gorda.di.uminho.pt/

EC contribution
1 218 643 €

start date
01/10/2004

duration
42