

# A Case for Relational GIS/GMA using Relaxed Consistency

Brian Coghlan, Trinity College Dublin, [coghlan@cs.tcd.ie](mailto:coghlan@cs.tcd.ie)

## Introduction

The draft paper on MDS2 that Ian Foster distributed to the Global Grid Forum (GGF) on the 6<sup>th</sup> April, 2001 [Czajkowski et al 2001], is very interesting. However the authors say in Section 4.1:

*“We believe that the join operation should not be provided as part of the basic query language. The utility of generalized joins is limited within the Grid due to our inability to view a consistent global state. Where a join operation is necessary, it can be supported via an optimized discovery service (see Section 5.2).”*

Clearly they are right regarding the inability to view a consistent global state, since any attempt to enforce global consistency with current database technology will lead to unacceptable inefficiency. I want here to erect a strawman argument to the effect that one should not abandon a fully relational approach just because of this, in the hope that this will stimulate further debate.

## Relaxed Consistency

In essence I argue for a relational model, but with relaxed consistency – in fact I propose that strict consistency is not only unnecessary, but that the lack of it is a healthy thing. The argument proceeds as follows:

- (1) Since efficient consistent distributed relational databases are an outstanding problem in computer science, we should relax consistency.
- (2) We may do this because true consistency on the grid is an illusion anyway.
- (3) The reason it is an illusion is that the authoritative source of information must always be the local resource, and hence any previously published information can never be fully trusted.
- (4) The reason this is so is because local policy must always prevail, regardless of previously published information.
- (5) As a result, any published information must be treated with scepticism, in that any decisions or actions resulting from it may fail when the resource is finally contacted. The information is an estimate rather than knowledge.
- (6) This doesn't mean that the decisions or actions are worthless, just that they are aspirations rather than directives.

Since this is so, any attempts to maintain global consistency are wasteful. It may be better to assume that inconsistency is there, and to assume that it is there for all the right reasons, not that it is a flaw.

A relational GIS/GMA need not be inefficient due to consistency issues.

## Other Issues

Contingent on this approach, there are several other points I would like to make:

- (7) Inconsistency may yield contradictions or livelocks. These will need to be specifically avoided.
- (8) It would be valuable to know the level of confidence (in a temporal sense) one could ascribe to information while forming an aspiration. How can one construct such a measure? There are at least three distinct cases: information about “continuous”, failure-prone, and unpredictable processes.
- (9) In the case of information from sources that would commonly be modelled as continuous

systems, one could use the statistical concept of chronological contiguity, perhaps by including something like the time-to-live (TTL) into the GGF grid timestamps [Gunter and Tierney 2001] as a loose indication of the expected temporal contiguity.

- (10) In the case of failure-prone resources, these aspects can be modelled as Poisson processes as is common in reliability studies.
- (11) In the case of unpredictable information, it may be possible to model this as a chaotic process.
- (12) If one has very little confidence in information, should it be discarded, and if so, by whom? Consider an archive: it is clear that a producer is not the best judge of the continued worth of the information [Fisher 2001]. This strongly suggests (for all cases) that filtering should always be the privilege of a consumer, not a raw consumer. Specialized filtering consumer-producers [Czajkowski et al 2001] would still fulfill a valuable role.
- (13) At present the grid resources needed by an application must be pre-specified, say in a resource specification language (RSL), and this gives rise to resource reservations. If we accept the above argument, however, these resource reservations become aspirations rather than directives.
- (14) Any such reservation may have alternate reservations created for invocation should the initial choice fail, or perhaps two or more may be acted upon in parallel and any redundant successes cancelled. These are typical stratagems. The essential point is that failure of a reservation, rather than success, has become the basic assumption.
- (15) Pre-specification limits the possibilities for using resources in an ad hoc manner. For example, if a CMS application discovers in mid-execution that it wishes to use an object that was not pre-specified, then the application may fail, for any number of reasons, such as localized conditions at the resource or more global conditions like authorizations. In particular it may fail because once these conditions occur, then at present the resource discovery is unlikely to be revisited.
- (16) Taken to the limit, alternate aspirations could be created for both pre-specified *and* spontaneous resource utilization, so that rejection by a resource results in fresh attempts rather than failure. In the case of the CMS example, resource discovery *would* be revisited, and in extremis the job could be suspended until the user rectified the situation. Hence, and atypically, pre-specification becomes an aid to efficiency, not a prescriptive bound.

## Summary

I have made a case for relaxed consistency in GIS/GMA. The argument applies equally well to a non-relational approach, but specifically benefits the relational model. It also has benefits for dependent services. I now invite debate.

## Acknowledgments

My thanks to Steve Fisher for his feedback.

## References

[Czajkowski et al 2001] Czajkowski, C., Fitzgerald, S., Foster, I., and Kesselman, C., *Grid Information Services for Distributed Resource Sharing*, DRAFT, distributed to [gis-wg@gridforum.org](mailto:gis-wg@gridforum.org) and [perf-wg@gridforum.org](mailto:perf-wg@gridforum.org), 6<sup>th</sup> April, 2001.

[Gunter and Tierney 2001] Gunter, D., and Tierney, B., *A Standard Timestamp for Grid Computing*, <http://www-didc.lbl.gov/GridPerf/papers/StandardTimestamp.pdf>, 22<sup>nd</sup> February, 2001.

[Fisher 2001] Fisher, S., *Relational Model for Information and Monitoring*, <http://www-didc.lbl.gov/GridPerf/papers/RelationalModel.pdf>, 16<sup>th</sup> February, 2001.