

E. Clévéde, G. Patau, J.P. Vilotte, D. Weissenbach

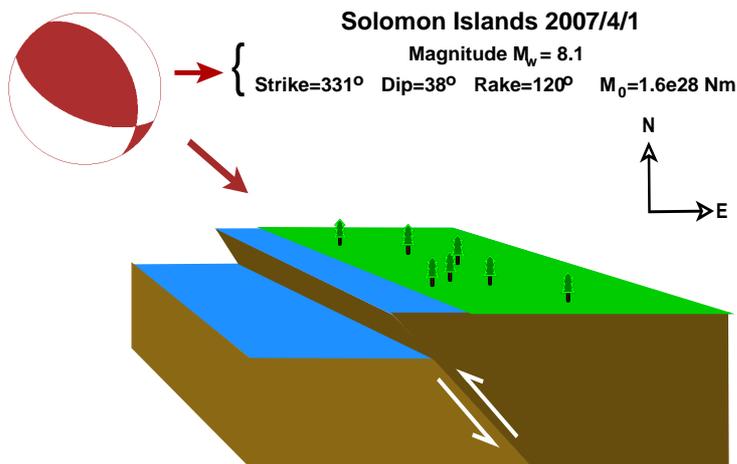
clevede@ipgp.jussieu.fr, patau@ipgp.jussieu.fr, vilotte@ipgp.jussieu.fr,
david.weissenbach@upmc.fr

EGEE'09

21-25 September, Barcelona

NA4 - VO ESR

An accurate description of an earthquake is given by the seismic Moment Tensor at the Centroid, the **CMT**.



In the case of a powerful earthquake ($M_w > 6.2$), embarrassingly parallel simulations lead to a solution (centroid location and moment tensor) determined by the best fit between simulated and observed data. It involves:

1. placing the quake's centroid at several (usually around 1600) locations paving a volume surrounding the suspected origin;
2. computing the resulting waves theoretically observed by the selected seismometers;
3. and then comparing them to the genuine seismic data produced by the event.

The usual process consists of a few runs needed to find significant fit variations between the tested locations by refining the selection of the referent seismic records.

On EGEE completion is obtained a few hours after the occurrence of the earthquake despite the necessary 500 to 2000 CPU hours.

Positions to inspect are dispatched by central server (Master) to running Grid jobs (Workers) during their execution¹.

Re-usable intermediate results (synthetic seismograms) are stored on the Grid and can save more than 80% CPU time for runs with refined data set or for aftershocks of this earthquake.

They can also be used in case of a next earthquake in the same area.

However, the EGEE File Catalog alone cannot allow to determine whether synthetic seismograms have already been computed for a given location.

Neither can it provide the list of the matching seismometers for which they were produced.

A recently introduced Grid accessible (by use of the GReLC² software) database service,

- updated on the fly by the Workers
- backing the Master process by displaying an accurate picture of the synthetic seismogram archive

allows systematic re-use of already computed synthetic seismograms.

As a side product, it is also possible to trace the use of the seismometers for the CMT determinations.

- The archive must now be consolidated and replicated, which also implies a distributed 2-tier organisation of the Resource Centers on which the application runs.
- The archive can be built as earthquakes are occurring, but also by pre-computing *synthetic seismograms* for determined areas.
- The growth of the archive is around 15GB per earthquake (estimated 300GB/year). A generated archive consisting of a loose paving of 4 of the Earth's main seismic regions would weight less than 5TB.

