

Distributed jobs on EGEE Grid infrastructure for an Earth science application: moment tensor computation at the centroid of an earthquake

Eric Clévéde · David Weissenbach · Benoît Gotab

Received: 17 September 2008 / Accepted: 29 April 2009 / Published online: 20 June 2009
© Springer-Verlag 2009

Abstract The Centroid Moment Tensor is a good summary of the fault rupture happening during an earthquake. We here describe a method for its determination in case of large earthquakes (with magnitude greater than 6.5) recorded world-widely by permanent digital seismic stations networks, as well as how deploying it on a Grid computing facility finally enabled its fast completion. Two different approaches of this process used over the last years are herein detailed and discussed. Some important impacts of this work are then highlighted and future trends are foreseen.

Keywords Centroid moment tensor · Earthquake · Embarrassingly parallel algorithm · Grid computing · Master/worker architecture · Synthetic seismogram

Introduction

The first quantitative information sought after an earthquake occurrence is, beyond the geographical location of its epicenter, the characterization of the event: what is the magnitude and duration of the earthquake, its depth and its focal mechanism, *i.e.* the fault orientation and the direction of slip along that fault. A rapid

estimation of these parameters is necessary before any further study. The seismic centroid moment tensor is an equivalent point source representation of these average characteristics.

Obtaining these information from seismograms is the mandatory first step of a seismic event analysis before further detailed investigations: ground deformation, kinematics of the rupture, seismicity analysis of a fault system, etc.

Also, seismological studies of the large scale structures of the Earth, from regional to global scales, are performed using a tomographic process in which thousands of synthetic seismograms are generated and compared to data. Catalogs of seismic moment tensors are used as input to generate these synthetic seismograms.

The centroid moment tensor is always determined for seismic event of scientific and/or social interest, but there are only two systematic catalogs available at global scale, including a rapid determination of the earthquakes source mechanism: the Global CMT project¹ and the USGS centroid moment tensor² solution. A few regional initiative also exist (e.g. by the USGS or the regional moment tensor catalog of the EMSC).³

Rapid computing of moment tensors needs a comprehensive strategy in which rapid real-time data access

Communicated by: H. A. Babaie

E. Clévéde (✉) · D. Weissenbach · B. Gotab
IPGP, 5 place Jussieu, 75005 Paris, France
e-mail: clevede@ipgp.jussieu.fr

D. Weissenbach
e-mail: david.weissenbach@upmc.fr

B. Gotab
e-mail: gotab@ipgp.jussieu.fr

¹<http://www.globalcmt.org/>

²United States Geological Survey, <http://www.earthquake.usgs.gov/>

³European-Mediterranean Seismological Centre, <http://www.emsc-csem.org/>

to a representative set of data, long computing times and a user friendly procedure are combined.

We first concentrated our efforts on the the core application of this process. This paper focuses on this application which consists in the computational procedure *per se*, and the software interface that takes care of the workflow management.

The computational procedure is itself divided into modules: data pre-processing, synthetic data computations, inversion procedure and results post-processing. Each of these elements can be modified in order to deal with different configurations: for example, depending on the type of data considered, any suitable method and/or Earth model can be plugged to compute the synthetic data.

As the method chosen is intrinsically distributed, deploying it on a Grid computing infrastructure that gathers several thousands of computers such as the EGEE⁴ project was seen as a solution to reduce the time spent on this operation.

The choice of a Grid computing environment has two main reasons. First, as the application is not a tightly coupled parallel one, it does not suits the usual requirements of high-performance computing facilities. Second, the possibility to interface the EGEE Grid environment to a database infrastructure, and to develop a user friendly interface, that could possibly be used by any interested person in a transparent way (through a web portal for instance), meet the requirements of a rapid seismic moment tensor determination procedure.

We present here the application in the framework of the determination of the centroid seismic tensor for large earthquakes, *i.e.* producing long-period seismic signal (above 100s period) on a global scale.

Operational constraints as well as other issues made this *gridification*⁵ an evolving process covered by Sections “CMT determination using EGEE computing Grid” and “The grEQ software”. It allowed to better understand the Grid and the application itself which had been underused before this opportunity because of its heavy computations needs. This experience lead to think of potential changes in the application structure, as described in the last paragraphs of this paper.

⁴<http://www.eu-egee.org/>

⁵The expression *gridification* refers both to the inclusion of a site into a Grid and to the adaptation of a software to derive benefits from a Grid (Briquet and Marneffe 2006). The latter meaning applies here.

Centroid moment tensor computation using a linear inversion method

Seismic source representation

The moment tensor

An efficient way to mathematically represent the average rupture process of a seismic event is the centroid moment tensor, which corresponds to a point source equivalent mechanism located at the barycenter (center of mass) of the seismic slip distribution, both in space and time (Backus and Mulcahy 1976a, b; Dziewonski et al. 1981).

Using this description, the computation of seismic wave excitation is a linear problem, making easier the determination of the source characteristics from the observations through an inverse problem.

The moment tensor $M_{ij}(i = 1, 3; j = 1, 3)$ is a system of forces described by a second-rank moment. The diagonal terms corresponds to linear dipoles, and the off-diagonal terms correspond to force couples. When assuming no net torque, the tensor is symmetric, reducing the number of independent components of the tensor to six. In the case of a source with no volume change, like a dislocation occurring along a fault, the trace of the moment tensor is zero, meaning no isotropic part, leading to five independent components. Isotropic part is however considered in the case of explosion or implosion, as, for example, in nuclear explosion, quarry blast or volcanic event.

Earthquake source mechanism interpretation

- Magnitude and duration: The size of an event is given by the seismic moment M_0 (expressed in $N.m$), which is directly obtained as it corresponds to the norm of the moment tensor, and its duration. The magnitude of a large event is commonly represented on the moment magnitude (M_w) scale (being an open scale) which is a function of M_0 only. M_0 (and thus M_w) is proportional to the area of the fault that ruptured and to the average amount of slip of this rupture.
- Source mechanism: There are different ways to interpret the moment tensor representation. In the case of earthquakes, the most common one is the decomposition into a best double-couple and a minimal compensated linear vector dipole (CLVD). The double-couple model is the preferred representation by equivalent body forces of a simple dislocation source. Its geometrical interpretation gives the fault plane orientation (strike and dip)

and the direction of the slip vector (rake) on this plane (Fig. 1). The double-couple gives actually two planes, one for each couple: the fault plane and the auxiliary plane, which can lead to ambiguity, but can be discriminated by the tectonic settings (Fig. 1).

The CLVD component can be interpreted in different ways, as complex rupture history or heterogeneity of the source medium (Julian et al. 1998).

Hence, the best double-couple and minimal CLVD representation assumes that the rupture process geometry is rather simple. In the seismic moment tensor determination procedure this representation is given as a by-product, but the seismic tensor itself allows different decompositions (as, for example, a composite rupture on different fault segments).

Seismic moment tensor determination procedure

The moment tensor is linearly related to the data when the centroid location in time and space is known. However, there is a trade-off between the centroid location and the mechanism when the moment tensor and the centroid are jointly inverted. Hence, in order to determine the centroid seismic tensor, we chose a mixed approach. The area around the first hypocentral determination is considered as a 3D parameter space to be explored, in terms of latitude, longitude and depth. The duration of the event is also explored as a 4th dimension of the parameter space. For each point of this exploration of the parameter space, a virtual source is determined by linear inversion, giving an associated

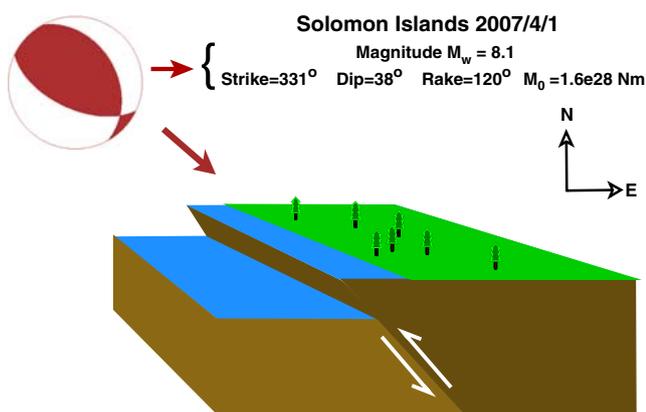


Fig. 1 Solution obtained for a Solomon Island Earthquake. The “beach ball” is the standard graphical representation of the moment tensor. The angles define the fault plane associated to the best double-couple. Here the fault plane is identified by the tectonic setting (subduction zone)

moment tensor. In this way, the global problem is divided into several small linear problems.

The linear inversion is performed in the complex spectral domain using relevant temporal windows of data and frequency range. The seismic moment tensor solution is determined by the best fit to the complex spectral data among all the moment tensors computed for each point of the parameter space.

Note that this approach is similar to the methodology used for real-time monitoring of a specific zone, as the GRiD-MT⁶ system in Japan, operating since 2003 (Tsuruoka et al. 2008).

Data selection and pre-processing

- Data: We use long-period three-component seismic data recorded worldwide. Usually data are extracted from the LH record channel corresponding to a sample rate of 1 Hz, over a time window of at least 5 h starting about 1 h before the event origin time. The data selection is first done on raw data to ensure that there is no instrumental problem (such as high level of noise, holes, glitches, signal saturation). A second selection is made after a coarse low-pass filtering, in order to avoid seismic traces with high long-period noise, and to obtain an azimuthal coverage as good as possible (the latter strongly affected by the uneven distributions of large seismic events and permanent seismic stations over the Earth).
- Input parameters:
 1. Time windows (from a few hundreds to a few thousands seconds, with 1 sample per second) are selected for each single trace, corresponding to the maximum signal. Several time windows can be used for each trace.
 2. A frequency range is determined for the spectral domain inversion, depending on the first estimation of the magnitude. Typically, for a large event of magnitude M_w 6.8 to 7.8, the frequency window is 4–8 MHz. These numbers can be modified for re-iteration of the inversion process if necessary.

Centroid moment tensor determination procedure

Latitude, longitude, depth and duration of the point source centroid being fixed,

1. 6 synthetic traces (seismograms) are computed by normal mode summation, for each data trace

⁶http://www.eri.u-tokyo.ac.jp/GRiD_MT/

corresponding to a given component (horizontal North–South, horizontal East–West, vertical trace) for a given station and the given source location. These 6 synthetic seismograms are computed using a virtual source consisting of one of the independent component of the tensor being equal to one, the others being equal to zero. The set of these synthetic seismograms are our so-called Green functions (note that even if 5 components of the moment tensor are actually inverted for—as we consider the non-isotropic case—the 6 Green functions are necessary in order to reconstruct synthetic seismograms to compare with data).

2. A first linear inversion, using the complex spectra of the selected time windows, is performed to obtain a preliminary solution.
3. Using the preliminary moment tensor solution, a synthetic seismogram is computed for each seismogram. A cross-correlation with the real trace is performed in order to determine a time centroid for each station.
4. Taking into account the local shift in time at each station, a final linear inversion is performed, giving a moment tensor solution for the fixed location.
5. The synthetic seismograms are constructed using the moment tensor solution and the value of the cost function for this point of the parameter space estimated.

CMT determination using EGEE computing Grid

The embarrassingly parallel⁷ (Foster 1995) nature of this CMT determination, each pretended source being completely unrelated to any other, makes it an ideal case of Grid application, and allowed it to be one of the first in the domain to ever be ported to a computing Grid.

Starting on the new EGEE infrastructure in late 2004, it was soon clear that a more elaborated design was needed to attain tremendous reduction of the duration of a run. Completing the several hundreds of required CPU hours—linear complexity regarding the number of traces in the input data set—was expected within half a day.

Nevertheless, despite its own young age and the small amount of committed resources (less than 10 Grid

clusters were allowing the application), the first results on EGEE were a success in some ways:

- starting from scratch, it allowed us to get familiar with the infrastructure while already using it for a genuine production;
- the solution was obtained noticeably faster, in about 2 days instead of a couple of weeks;
- the lessons learnt could lead to the first objective: divide the execution time by a factor of several hundreds.

In this section, we detail difficulties we faced during this period, and how significantly they impacted the process.

Simple adaptation on the Grid

During its first year on the Grid, the application was run as follows: there was one separate job for every triplet composing the parameter space. Each job including the software itself (statically compiled executables and shell scripts) and the selected seismograms, was submitted to the Grid through a single Resource Broker (the service in charge of dispatching the jobs to entrusted clusters of the Grid's Resource Centers).

Jobs had first to retrieve the 60 MB Earth model from specialized Grid Storage Elements where it had been once recorded. Upon completion (after an average of 30 min) their results were sent back to the Resource Broker before being collected back.

This almost canonical *gridification* tried to reproduce the way it could have looked if submitted to locally available managed—and eventually volatile—resources, with full trust in the Resource Manager. However, EGEE is a distributed Grid and despite some luck with several factors (such as limited volume transfers allowing the use of the Grid's Sandboxes feature), and a bit less with others like job durations, this apparent simplicity turned out to be quite complex to handle.

As already said, the full completion usually needed about 2 days, compared to two weeks when relying only on local computing facilities. The main difference was that it did not need any human input for the latter, while the former one required a person to pilot the application.

Facing the Grid

Depending on many factors, some jobs could fail, while others would remain endlessly queued in the batch systems, when successful ones released after half an hour the CPU they had hardly obtained to any other Grid job.

⁷*Embarrassingly parallel* means that the computation consists of a number of tasks that can execute more or less independently, without communication.

Fig. 2 Client pseudo-code

1. Get static data and software from Storage Element; in case of failure go to **step 3.** ; set status to joining
2. Send status (*joining, done or leaving*) to server. If “got coordinates triplet back” do:
 - i. If existing, get archived synthetic seismograms for coordinates triplet
 - ii. For each trace in data set:
 - (a) Compute non existing synthetic seismograms
 - (b) Compare synthetic seismograms against it (Linear inversion process)
 - iii. Update synthetic seismograms archive
 - iv. Estimate run time from **step (i.)** to now. Compute mean run time (MRT)
 - v. Save results back to Storage Element
 - vi. If “more available CPU time than MRT” set status to *done* ; go to **step 2**
Else set status to *leaving* ; go to **step 3**
3. Terminate gracefully

One had to frequently resubmit failed jobs, to cancel and immediately resubmit those for which it seemed improbable that they would quickly be run. Without this human input, the CMT determination remained unsecure.

This task was later a bit assisted, but it still needed a Grid expert to be able to determine or to guess which jobs to eventually redirect to a site where they could be executed sooner.

The case of a small but usual initial investigation area of square of 1° on ground (with a 0.1° step) and 25 km deep (with a 2.5 km step), resulting in a $11 \times 11 \times 11$ cell space and thus globally needing about 665 CPU hours for the 1331 jobs; their submission through a single Resource Broker lasts for about two hours. After this time, some jobs will already have completed and others will still be running, but on bad cases hundreds of them could still remain non started.

As a first remediation, grouping several computations within the same job was adopted but although more scalable. It did not actually speed up the process because of the reliability issue; a failing job, if not detected soon enough, would longer delay the end of the process due to its increased duration. Unmonitored (like overnight) runs were still bound to partly fail.

The grEQ software

The figured improvement is to get rid of the dependency induced by the statical assignment of the varying parameter (the position of the source location) to unpredictable jobs. The idea is to delay this binding until

it is clear that all chances of success are met. Grid jobs are turned into clients that repeatedly request tasks to carry out, but only if fully capable. This allows to delegate the responsibility of monitoring the process to an automated server. Interacting with the Grid, although simplified by the server which also generates the jobs, is still possible for the human operator.

Implementing the method

By the end of summer 2006, the introduction of grEQ, a new home-designed software implementing a master/slave architecture, finally ensured CMT determination. Grid jobs and tensor computations are now monitored by a server which main occupation is to dynamically distribute the triplets covering the parameter space to its clients. These compute the synthetic seismograms for the location and perform the inversion of the earthquake records, save their result, and return to the server a request for another triplet, or notifying the exhaustion of their granted execution time.

Client description

The client, shown in Fig. 2, is a very basic rewriting of the previously used CMT shell wrapper (boxed items on Fig. 2), wrapped itself within the loop of a python script that is in charge of all initiations (and afterwards terminations) and newly introduced communications.⁸

⁸Steps (i.) and (iv.) described in “[Changing the nature of the computations](#)”.

The data shipped with the job `InputSandbox` still includes a subset of the recorded seismograms resulting from the earthquake. It also contains a unique code generated by the server and used for integrity checks that identifies the client among all.

Server actions

The server:

1. authenticates the connecting clients implementing the GSI⁹ (Foster et al. 1998) authentication methods;
2. ensures they are in possession of the intended data set (check that the client key is known);
3. check their request (one of: *joining*, *done*, *leaving*);
4. eventually accounts client and decides the attribution of a source location in case of a non *leaving* request.

The attributed location may be one that had already been assigned to another client, in case this client is supposed to be missing after not recontacting the server on time (some delay is tolerated) while compared to its records or to the slowest client. A client may be given a termination order in case all positions are done or attributed and still expected, or the client has revealed to slow. In the general case where an unattributed location is assigned, it is randomly selected.

When a client request is *done* or *leaving*, the server may launch a background process to upload the results for the last position this client was attributed from the Storage Element.

A client claiming to be a new one (request is *joining*), but for which the server already possesses accounting data, reveals that for some reason it has been restarted. The previously assigned position is supposed not to have completed and is immediately re-attributed to this same client.

Immediate outcomes

Using this software, it is now usual to effortlessly determine the centroid and its associated tensor a few hours after the data are preprocessed. Configuring the server—describing the parameter space—, starting it and submitting the jobs through the `gLite` middleware (EGEE 1 2004) is also faster than using the legacy DataGrid Resource Broker. The server preparing the required additional files for the client submission and

performing the file transfers, the user-friendliness of the application also improved, despite the new need to collect output from Storage Elements while it used to be transferred by the `OutputSandbox`.

Also, inspecting the randomly positioned first computed tensors is found to be a sufficient quality control when determining whether data pre-processing and selection need to be modified, neglecting the need of preliminary “calibration runs”.

CMT determination typical work flow

A typical work flow is shown in Fig. 3: an alert is raised by the USGS (or from another organization devoted to the same task) when an earthquake occurs. GEOSCOPE¹⁰ prepares in the next hours a package composed of all available recordings for the event and notifies its URL by e-mail (any other equivalent data from other networks can be used).

The operator then selects from this SEED¹¹ (SEED 2009) volume the seismic traces depending mainly on their quality, but also on the distance to the earthquake location. The data pre-processing is quasi automated, leaving quality control to the operator.

This input data set is then transferred to the `grEQ` server host which happens to also be a `gLite` User Interface (*i.e.* equipped with the compulsory middleware for job and file management on EGEE). Finally, after instantiation of the server the user submits the created jobs to the Grid.

Synthetic Green functions are calculated in the process as needed by parameter space cells. They are automatically stored in order to avoid re-computation in case of results improvement triggered by unsatisfying outcome—dash arrows on Fig. 3—or for a probable aftershock.

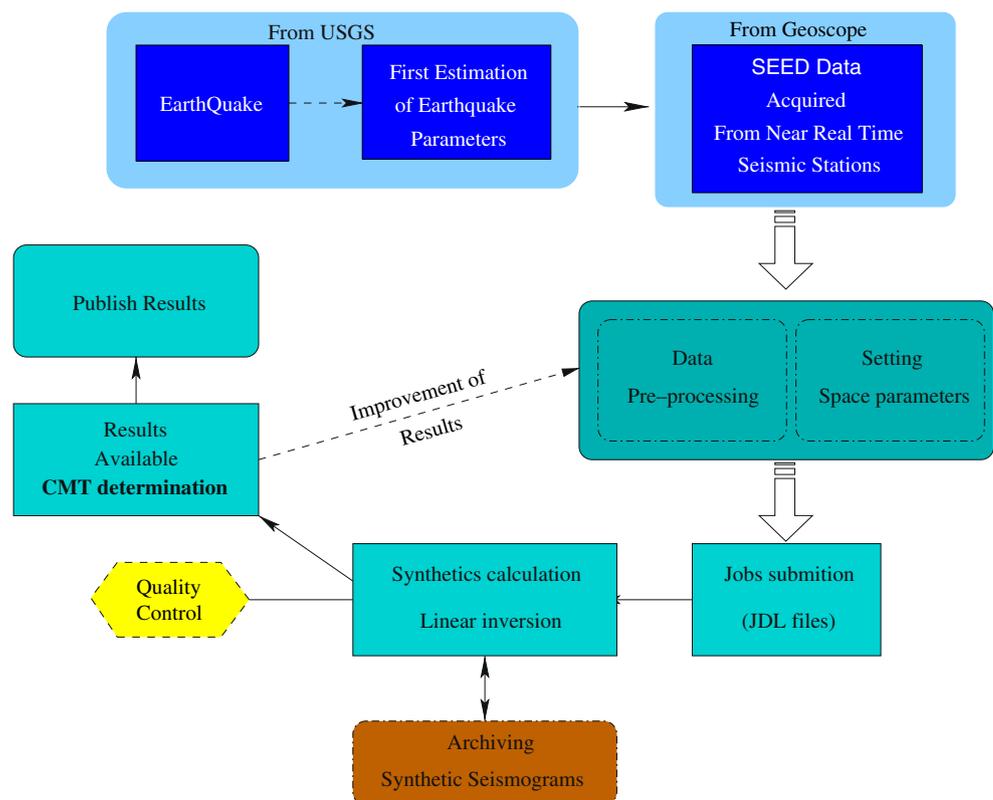
Possibly a few other runs will still be necessary before a fine solution is obtained, but only triggered by physical parameter changes such as spatial grid shift or changes in the input data. Here, available additional program features such as the possibility to operate

⁹Grid Security Infrastructure.

¹⁰The GEOSCOPE network is composed of 32 wide band seismic stations covering the globe. GEOSCOPE is the French component of the networks of seismic monitoring know as “very broad band”. It network takes part in localizing of earthquakes on the whole planet, in determining the rupture mode of the faults causing them and also in a true three-dimensional radiography (tomography) of the interior of our planet. GEOSCOPE is also associated to the development of international monitoring systems, in particular for tsunami related risks in the Indian Ocean and the Antilles.

¹¹Standard for the exchange of earthquake data.

Fig. 3 Work flow used for the determination of the seismic moment tensor at the centroid of an earthquake



several servers at once are also helpful. Finally, Final CMT solutions are published on the GEOSCOPE website.¹²

Impacts of grEQ

After drastically reducing the execution time from weeks to hours, and ensuring the reliable execution of the whole processing, thanks to the two stages of *gridifications* we described, the application is nowadays used on a efficient routine base.

We address in this section some considerations on the job management efficiency obtained by grEQ on the EGEE Grid, and the experience of using our application as a test case for Grid project deployments.

Grid considerations

Usually about half of the submitted jobs manage to become active clients during the run. Late clients stand for a last chance to complete the picture in case the final computations remain pending. It was considered the best policy to explicitly submit jobs to each and

every Resource Center in numbers matching a fraction of their size so as to be able to take the opportunity of undetected available execution slots and to avoid job attraction by some Grid “black hole”.

However, despite these overheads, the current architecture is lighter when weighted on the Grid’s balance:

- fewer jobs (maybe 600 for the previous example with 11^3 computations and 900 for a larger one with over 2500 cells);
- no more massive job cancelings—a long known and yet unsolved Resource Broker weakness—and resubmissions;
- fewer file transfers as the same static data will be used for several spatial positions compared to the former one to one ratio; even with synthetic seismograms upload, the displaced volume decreased.

grEQ is not aimed at completely hiding the Grid to the user, but at proposing an easy-to-use specialized CMT determination tool in the scope of a generalist distributed computing service. As for instance, a standard overnight run or one of an unusual kind are not handled the same way, the global position/job ratio and some more flexibility opportunities are left to the user. However he/she has the possibility to submit additional clients to sites executing all the ones received in the first place, so as to try to increase the rate of moment tensor

¹²http://geoscope.ipgp.jussieu.fr/CMT/Default_cmt.htm

computing. Although the application will reliably allow a Grid novice or a careless user to complete a run, a Grid expert will still have ways to deploy his talents while using grEQ.

The latter introduced synthetic seismogram archival (described in the next section) has important outcomes in case of strong aftershocks (as seen in January 2009 as from the 4 major quakes, one had one aftershock and another two), but is most useful for every frequent situation where some changes are needed in the input, as time invested in every run can be effectively restored to the following ones.

CMT as testing platform

During the last two years, the CMT application was integrated within the surveys of the DEGREE¹³ project (Som de Cerff et al. 2009; Tran et al. 2009), an European Union Specific Support Action aiming to promote and widen the use of the Grid technology towards the Earth Science (ES) community and to create a bridge with Grid communities. Indeed, the application as well as other ones from multiple ES fields of research formed the backbone of the DEGREE test suite. The objectives of this test suite is to address ES requirements from actual applications to Grid developers; this allows to monitor Grid middleware, tools, and finally to get an overview on how the requests are fulfilled.

The embarrassingly parallel aspect of the application, the fast turnaround needed, and the ease of deploying it, made CMT computing an good candidate to be part of the DEGREE test suite; also noticeable was the fact that Grid use cases and performance measurements were already existing for the application. This test suite could allow developers to improve their tools and/or middleware to be able to respond to requirement from ES community. In this frame, using the french research Grid Grid5000 the CMT application featured as a third party program to the DIET¹⁴ tool.

The EGEE project also took notice of some advices to improve the way applications expecting fast response from the Grid would meet their needs. Several solutions like job prioritization or “on alert” jobs were studied and one named *short deadline jobs* (Germain-Renaud 2006) (or also *sdj*) in the EGEE terminology for short jobs that need immediate execution was finally implemented.

Unfortunately, the trade-off for the low scheduling is that the associated queues usually have very low limits on the CPU time, and allow the execution of these jobs on already busy processors. The amount of CPU involved for the computing of Green functions thus prevents them from being included within *short deadline jobs*, without completely neglecting the usefulness of this feature in the scope of our application, as explained in “[Incubation of next generation software](#)”.

New developments

In the meanwhile, improvements of the Grid’s file catalogs, moving from the legacy DataGrid Replica Location Service (Chervenak et al. 2004) to the LCG File Catalog (Munro et al. 2006), and the increasing amount of storage resources offered a good opportunity to save CPU time when several runs are needed before a fine solution is obtained. Moreover, their availability before the occurrence of an earthquake would allow an almost immediate CMT determination.

Changing the nature of the computations

Computing Green functions accounts for 80% of the computations, so synthetic seismograms are now carefully recorded and available for reprocessing of the same earthquake or for the processing of another event in the same area (e.g. an aftershock). Only possibly missing seismograms (needed if the input data contains new recording stations) are then computed and added to the archived set.

One side effect is that CMT computations used to have a fixed duration depending only on the computing power available on the execution node, and it now may vary with location. If the operator shifts the bounds of the explored area, some cells will have no pre-determined Green functions while others will. In such a case, the software has to deal with jobs sometimes lasting much longer than expected; usually considered as crashed while they are correctly performing, the same computation can be requested from other clients before the first one can complete it.

To overcome this issue, the estimation of the next computation duration is left to the client once that it can precisely figure it, after having downloaded and ensured the integrity of the seismograms.

Incubation of next generation software

The next generation of the CMT determination application could consist in just inverting data from the event

¹³Dissemination and exploitation of GRIDs in Earth science.

¹⁴Distributed interactive engineering toolbox, <http://graa.ens-lyon.fr/~diet/>

records using synthetic seismograms uploaded from an indexed archive stored on some data specialized facility—short deadline jobs can be used here—after having been systematically computed (as it is done by real-time determination operation in specifically monitored regions, e.g. Tsuruoka et al. 2008).

This would allow either triggered, knowing that an earthquake just occurred with explicit event traces, or regularly scheduled inversions as some seismic data can also be acquired in near real-time.

Synthetic seismograms computed in the course of earthquake analysis could be a starting point with the covered volume expanded as storage becomes available and selected areas are added.

The volume of the intended synthetic seismogram archive and catalog would incite to implement it on top of a distributed architecture, while computations would probably take place only in a few sites close to the archive locations, as moving such volumes along distances should be avoided when possible.

Possible implementation

Our design of such a software could rely on EGEE on which computations can be carried out and the archive maintained by a low intensity but continuous process, including as well those created during real earthquake studies. It still needs to be proved that the high amount of mobilized storage meets its value if used this way.

One could expect about 10Mb in the first times for the volume of compressed synthetic seismograms for one coordinate triplet. This is enough to store the vertical components for the 20 most frequently used stations, and the horizontal components for the 5 stations that are less impacted by seismic noise and thus the most likely to be used because of the better signal quality.

In a terabyte, 100,000 cells can then be stored, representing a volume of only 50 km deep over an area of $10^\circ \times 5^\circ$ (with a 0.1 deg step in lateral discretization, and a 2.5 km step in depth), which would cover, for instance, only portion of the Indonesian or Japanese regions.

The EGEE data model would be to maintain the archive in many different locations, but 3 to 5 strategically chosen ones, as well as a more careful Resource Center selection to avoid long distance transfers, should be enough considering the size of the application: it will still find close enough execution slots to be able to run hundreds of jobs simultaneously.

The database system, although securely readable and writable (for live updates) by Grid jobs, will probably remain centralized, in case of improbable—despite the

additional intermediate middleware layer needed between a standard open source SQL database and Grid authentication and authorization standards—scaling issues. In our case, the database interface will be composed by a GRelC¹⁵ Data Access Service service (Fiore et al. 2008) hosted at IPGP's EGEE Resource Center, on top of a postgresql database which will also be available for statistics concerning the GEOSCOPE network usage.

Validating the model

Until recently, the few gigabytes generated with an earthquake were kept no more than about a year. The process of indexing them in a database has already started. The next step consists in the concentration and the replication of this distributed archive. Then, by concentrating efforts on a few active seismic zones, a first significant step will soon be completed.

As the computing power will still be at hand, for a start the archive will be using a more coarse grained 3D-grid, so covering a much larger volume so as to raise the chances of using it in the short term. After or in parallel with a quick archive-based determination of the most probable centroid location, a finer grained study will be performed with grEQ to confirm the precise location, and additional Green functions recorded.

Conclusion

The CMT determination application took many advantages from its porting to the EGEE infrastructure.

First, the Grid facilities, together with the implementation of the grEQ software, allows fast determination of large earthquakes mechanism on a routine base and in a reliable way. The results obtained now feature as a product of the GEOSCOPE network.

Moreover, this *gridification* helped the application in particular, and as Earth Sciences applications in general, to reach a wider audience in the Grid communities through presentations and demonstrations, as well as it also helped to depict some of the Grid technologies to scientists needing flexible and on-demand large amount of CPU time in a collaborative framework.

As it was stated in the introduction, rapid computing of moment tensors needs a comprehensive strategy in which rapid real-time data access to a representative set of data, long computing times and a user friendly procedure are combined.

¹⁵Grid relational catalog.

The core application of this procedure being now efficiently implemented, the other steps can be considered. Three main objectives are scheduled for the future developments. The first one is to perform automatic data pre-processing and results post-processing. The second one is to consider the possibility of an interface between the application and a seismic database infrastructure. The third one is to develop a user friendly interface, such as a web portal, in order to open the application to a wider community. One lead is the effort made in databases access and the development of tools for improved service and data analysis in the framework of the NERIES¹⁶ European project. To fulfill these objectives will lead to a complete achievement of rapid seismic moment tensor determination procedure.

Acknowledgements This work makes use of results produced by the Enabling Grids for E-sciencE project, a project cofunded by the European Commission (under contract number INFSO-RI-031688) through the Sixth Framework Programme. EGEE brings together 91 partners in 32 countries to provide a seamless Grid infrastructure available to the European research community 24 h a day. Full information is available at <http://www.eu-egee.org>. This is IPGP contribution no. 2523.

The authors wish to thank Geneviève Patau. Without her efforts and patience, precise CMT of many major earthquakes would still remain undetermined. The authors also thank Jean-Pierre Vilotte and Alexandre Fournier for their useful comments.

References

- Backus G, Mulcahy M (1976a) Moment tensors and other phenomenological descriptions of seismic sources, I, continuous displacements. *Geophys J R Astron Soc* 46:341–361
- Backus G, Mulcahy M (1976b) Moment tensors and other phenomenological descriptions of seismic sources, II, discontinuous displacements. *Geophys J R Astron Soc* 47:301–329
- Briquet C, de Marneffe P-A (2006) What is the Grid? Tentative definitions beyond resource coordination. Technical report, University of Liège, Belgium. http://www.montefiore.ulg.ac.be/~briquet/what_is_the_grid_2006.pdf. Accessed 2 Feb 2009
- Chervenak AL, Palavalli N, Bharathi S, Kesselman C, Schwartzkopf R (2004) Performance and scalability of a replica location service. In: Proceedings of the international IEEE symposium on high performance distributed computing (HPDC-13). <http://www.globus.org/alliance/publications/papers/chervenakhpdc13.pdf>. Accessed 30 Jan 2009
- Dziewonski AM, Chou T-A, Woodhouse JH (1981) Determination of earthquake source parameters from waveform data for studies of global and regional seismicity. *J Geophys Res* 86:2825–2852
- EGEE Middleware Architecture, EGEE EU deliverable (2004) DJRA1.1. <https://edms.cern.ch/file/476451/1.0/architecture.pdf>. Accessed 30 Jan 2009
- Fiore S, Cafaro M, Mirto M, Vadacca S, Negro A, Aloisio G (2008) The GRelC project: state of the art and future directions. In: Grandinetti L (ed) High performance computing and grids in action, vol 16. IOS, Amsterdam, pp 331–334
- Foster I (1995) Designing and building parallel programs, section 1.4.4. Addison Wesley, Reading. <http://www-unix.mcs.anal.gov/dbpp/text/Node10.html>. Accessed 30 Jan 2009
- Foster I, Kesselman C, Tsudik G, Tuecke S (1998) A security architecture for computational Grids. In: Proc 5th ACM conference on computer and communications security conference, pp 83–92. <ftp://ftp.globus.org/pub/globus/papers/security.pdf>. Accessed 30 Jan 2009
- Germain-Renaud C (2006) Short deadline jobs Working Group first report (2006). <http://egee-intranet.web.cern.ch/egee-intranet/NA1/TCG/wgs/sdj.htm>. Accessed 30 Jan 2009
- Julian BR, Miller AD, Foulger GR (1998) Non-double-couple earthquakes, 1- theory. *Rev Geophys* 36:525–549
- Munro C, Koblitz B, Santos N, Khan A (2006) Performance comparison of the lcg2 and glite file catalogs. *IEEE Trans Nucl Sci* 53(4/2):2228–2232
- Standard for the Exchange of Earthquake Data (SEED) (2009) Manual Version 2.4. Edited by the international federation of digital seismograph networks (FDSN)
- Som de Cerff WJ, Petitdidier M, Kraus J, Plevier C, Schwichtenberg H, Tran V (2009) Contribution to test suite and specification report, DEGREE EU deliverable 1.3. http://www.eu-degree.eu/DEGREE/internal-section/wp1/DEGREE-WP1-D1_3-Test-Suite-Organization-v1.5.pdf. Accessed 30 Jan 2009
- Tran V, Habala O, Lonjaret M, Poliakov A, Schwichtenberg H (2009) Contribution to test suite and specification report (job management and control), DEGREE EU deliverable D3.2. http://www.eu-degree.eu/DEGREE/internal-section/wp3/DEGREE-WP3-D3_2-Test-Suite-1.1.pdf. Accessed 30 Jan 2009
- Tusruoka H, Kawakatsu H, Urabe T (2008) GRiD MT (Grid-based realtime determination of moment tensor) monitoring the long-period seismic wavefield. *Phys Earth Planet Int* (in press)

¹⁶Network of Research Infrastructures for European Seismology, <http://www.neries-eu.org/>