

Data Handling of a High Resolution Ground Motion Simulation

Authors: Geoffrey Ely(SIO), Qiao Xin(SDSC), Marcio Faerman(SDSC), George Kremenek (SDSC), Boris Shkoller(SIO), Steven Day(SDSU), Kim Olsen(UCSB), Bernard Minster(SIO), Reagan Moore(SDSC)

Abstract:

A Pathway 3 Data Generation and Handling Exercise was performed to demonstrate the data processing lifecycle for a Rupture Dynamic Model program.

Geoffrey Ely (gely@ucsd.edu) Marcio Faerman (faerman@sdsc.edu)

Introduction:

Large scale ground motion simulations, capable of representing frequencies of engineering interest ($> 1\text{Hz}$) for models of geological interest ($> 200\text{Km}$), are extremely demanding computationally. Furthermore, such high resolution runs may produce data sets that are multiple terabytes in size. Due to the high computational cost involved in running these simulations, it becomes critical to archive the produced datasets making them accessible for future use without the need of re-computation. Users should be able to discover, query, retrieve and analyze the archived data. Further, a carefully selected collection of such data sets will be used as benchmarks for assessing accuracy and performance of future simulations, developing post-processing software such as visualization tools, and testing data and metadata management strategies.

Data handling infrastructures called Digital Libraries provide such capabilities. Digital libraries organize datasets in collection hierarchies, associate metadata with data entities, provide querying capabilities, data transformation services, Internet publication and data preservation including operations such as, replication, access control, auditing and authentication.

We have run a high resolution ground motion simulation with the DFM rupture dynamics and wave propagation code and organization of approximately 10 Terabytes of output data and 1.5 million files, using the digital library technology provided by the SDSC Storage Resource Broker.

The DFM (Dynamic Fault Model) models spontaneous rupture on a planar fault within a 3D isotropic viscoelastic solid. The viscoelastic equations of motion are approximated with second order finite-differences. The fault is modeled by applying a slip-weakening, frictional, internal boundary condition. Parallelization is accomplished with domain decomposition and message passing using MPI.

DATA COLLECTIONS:

The SDSC Storage Resource Broker (SRB) provides the data handling capabilities to manage the Terabyte scale simulation output. The SRB provides support for ingestion, organization, description, preservation and access of datasets. The generated dataset is stored in the SRB collection:

`/home/mfaerman2.sdsc/SCEC-CME/Simulations/DFM`

which can be accessed from a web browser at

<http://www.sdsc.edu/SCEC/>, or with any other SRB client.

SUB-COLLECTION HIERARCHY: Not only the output data was stored. Other sub-collections containing information regarding the simulation are organized in the following way:

Input - Input metadata and input files.

Program - Source code metadata and source code files

Run - Simulation run metadata and run log files

Output was further subdivided as following:

Output/Velocity - Velocity vector field data

Output/Stress - Stress data

Output/vis_files - Velocity magnitude data used in Steve Cutchin visualization

METADATA:

Metadata were associated with each file of the simulation collections. The metadata attributes include descriptive information about the simulation run, simulation input parameters, the computational infrastructure, the physical geometry of the problem, and output structure. Below we include an example of the metadata associated to the simulation sub-collections.

Metadata for SRB collection /home/mfaerman2.sdsc/SCEC-CME/Simulations/DFM

Metadata describing entity DFM

DC.description : 4D wave propagation data

DC.subject.scheme : LCC

DC.subject.content : QE521-545

DC.type.scheme : DCMIType

DC.type.content : Dataset

DC.rights.Lang : en

DC.date.created : Thu Aug 28 09:09:11 PDT 2003

DC.creator.Name : Geoffrey Ely

DC.creator.Email : gely@ucsd.edu

DC.contributor : Qiao Xin, Marcio Faerman

DC.publisher.Name : Marcio Faerman

DC.publisher.Email : mfaerman@sdsc.edu

Metadata describing entity Input

PhysicalInputs.CoordinateSystem : Cartesian XYZ

PhysicalInputs.Volume.MinX : 0

PhysicalInputs.Volume.MaxX : 2.807

PhysicalInputs.Volume.MinY : 0

PhysicalInputs.Volume.MaxY : 5.607

PhysicalInputs.Volume.MinZ : 0

PhysicalInputs.Volume.MaxZ : 5.607

PhysicalInputs.TimeSeries.EndTime : 0.15

NumericalInputs.TimeSeries.NumberOfSteps : 3000

NumericalInputs.TimeSeries.TimeIncrement : 5e-05

NumericalInputs.3DData.NumPointsXdirection : 401
NumericalInputs.3DData.NumPointsYdirection : 801
NumericalInputs.3DData.NumPointsZdirection : 801
NumericalInputs.3DData.GridIncrement.Xdirection : 0.007
NumericalInputs.3DData.GridIncrement.Ydirection : 0.007
NumericalInputs.3DData.GridIncrement.Zdirection : 0.007
NumericalInputs.FreeSurfaceBoundaryCondition : natural

Metadata describing entity Program

ProgramInfo.Name : dfm
ProgramInfo.Description : DFM: Dynamic Fault Model
ProgramInfo.Version.User : Thu Aug 28 09:09:11 PDT 2003
ProgramInfo.Authors : Geoffrey Ely, Steve Day, Boris Shkoller
ProgramInfo.Affiliation : UCSD, SDSU
ProgramInfo.ProgrammingLanguages : C, Fortran
ProgramInfo.SystemSoftware : MPICH, bash, gawk, gmake, perl
ProgramInfo.Version.User : Thu Aug 28 09:09:11 PDT 2003

Metadata describing entity Run

RunID.Description : High resolution DFM foam rubber simulation
RunID.Submitter.Name : Qiao Xin
RunID.Investigator.Name : Geoffrey Ely
RunID.Date.Submitted : Thu Aug 28 09:09:11 PDT 2003
SystemID.HostName : tf005i
SystemID.OS.Name : AIX
SystemID.OS.Version : 1
SystemID.OS.Description : AIX tf005i 1 5 006025764C00, SDSC Blue Horizon
SystemID.NumberOfProcessors : 512

Metadata describing entity Output

OutputDescription.OutputFormat : 4 byte floats
OutputDescription.StorageFormat : 4 byte floats
OutputDescription.Endian : big

Metadata for collection /home/mfaerman2.sdsc/SCEC-CME/Simulations/DFM/Output/Velocity/02854

Metadata describing Time Step 2854

TimeStep : 2854

QUERY CAPABILITIES:

An user can query the dataset using the MySRB "Browse Query" button. Metadata attributes appear at a pull-down bar and the user can discover which files match the specified queried attributes. This feature is extremely useful with such a large data collection, ranging 1.5 million files.

VISUALIZATION:

Steve Cutchin has been able to render a high resolution visualization for the velocity magnitudes generated in this simulation. The large dataset is rendered efficiently since the processing is performed using multiple processors in parallel.

In a couple of months Steve Cutchin's rendering tools should have the following capabilities:

- Renderer configuration using a web page (web-service)
- Data paged and pipelined straight from the SRB into the parallel renderer
- User able to specify interactively which SRB collection she wants to be rendered
- Flexibility to plugin different camera files

Current tools and info are available on the web under: <http://vistools.npaci.edu/>

RDF Description:

In order to test our capabilities for handling large data sets, we performed a rupture dynamics simulation that generated 10 terabytes of output data. The model size is 400x800x800 nodes integrated over 3000 time steps. We save three components of velocity at every node and every time step as well as six components of stress at every node and every 10th time step.

The simulation code ran on NPACI Blue Horizon. Using 512 processors the code ran for 12 hours for a total of approximately 6000 CPU hours. Internal code timing information showed that approximately half of the compute time was spent performing i/o and even this was achieved only after significant i/o optimizations were made to the code. The style of i/o used for this simulation saves a separated file for each processor. This scheme results in fast i/o at run-time, but requires expensive post-processing to reassemble the separate files into a single grid. For this type of simulation where such a large amount of output is saved, the total run-time (execution time plus any required post-processing) is dominated by i/o (or "i/o bound"). This prompted us to incorporate parallel i/o into our simulation code. Parallel i/o allows saving a single file from multiple processors. Parallel i/o is slower due to the need to synchronize and coordinate across the processors, but the performance penalty is much less than that of post-processing. Parallel i/o will be used for all future i/o bound simulations.

The simulation was loaded into the SDSC Storage Resource Broker over a time period of about 20 days. This occurred during a period of hardware reconfiguration resulting slower than normal performance. A more typical time estimate for loading 10TBs is five days and will soon be less than 24 hours when planned improvements by SDSC are in place. Clearly, good network performance is critical for handling large data sets. The SDSC metadata handling facilities were used to organize and document the simulation.