SINBAD Consortium 2010

Software releases and architecture

Henryk Modzelewski

SLIM
University of British Columbia

SLIM

Outline

- Challenges
- Development environment
- Knowledge transfer
- Past software releases
- Next software release
- Conclusions

Challenges

- Seismic applications are:
  - memory demanding
  - computationally intensive
  - IO intensive
  - we use large-scale-optimization iterative algorithms
- SLIM's technology is driven by multi-d (d>2) non-separable transforms
- Academic environment calls for
  - code reuse & reproducibility
  - fast learning curves

Programming tools

- In-core computations:
  - **Matlab** and **Java**
  - **C/C++/FORTRAN** for low-level implementations
- Seismic tools for data input/output:
  - Seismic Unix, Delphi, **SegyMAT**, **Madagascar (RSF)**
- Scripting tools:
  - Unix shells, **SCons**, Python
- Parallelism
  - **MATLAB** and **Java** (for parallel IO)

Hardware

- Desktops (for small-problem prototyping)
- Cluster (for parallel development and computations
  - 288 CPU cores (Intel 2.66 GHz)
  - 16 TB of disk space
  - Funding:
    - contribution to initial purchase (2006)
    - departmental effort to upgrade (2009)

Development in MATLAB

- Students typically have already experience with MATLAB programming
- MATLAB introduced the capabilities that allow more efficient algorithm development:
  - support for parallel computations
  - support for distributed arrays/vectors
  - MATLAB interfaces with other low-level languages
- Our goal is to isolate students , as much as possible, from low-level programing

Development in MATLAB

- MATLAB addresses our major challenges:
  - massive parallel computations
  - capability to handle large memory footprint
  - faster learning curve (less software to learn)
- Unfortunately MATLAB is not free
- But there is vast amount of scientific software developed in MATLAB

Development in MATLAB

- There is still need to built or utilize efficient low-level routines and interface them with MATLAB, like:
  - 3D parallel curvelet transform
- Some of the existing MATLAB code is being converted or extended to allow parallel computations, like:
  - SPOT
  - SPGL1
- There is still need to implement utilities for out-of-core processing and data input/output

Development in MATLAB

- Current software projects:
  - SPOT extensions for parallel processing using pSPOT
  - pSPOT: parallel extensions to SPOT
  - both are subjects of the next presentations

Development in MATLAB

- Future projects:
  - further development of pSPOT
  - improvements to interfacing out-of-core data
    - investigating JavaSeis (www.javaseis.org)
    - improvements to RSF-MATLAB interface
  - MATLAB implementation of parallel 3D curvelets
  - nSPOT: non-linear extensions to SPOT

Programming staff (future)

- We intend to maintain 4 positions to accommodate
  - part-time senior programmer
  - 2 full-time junior programmers
  - part-time system administrator
- Focus on supporting MATLAB development

Knowledge transfer

- Publications
- Software releases
  - algorithms
  - applications' demos
- Webinars (monthly/bi-monthly)
  - presentations
  - software demonstrations
  - software tutorials

Past software releases

- Software Highlights:
  - https://wave.eos.ubc.ca/Software/SINBAD/highlights/
- Software Releases:
  - https://wave.eos.ubc.ca/Software/SINBAD/

- Accessible only to consortium members
  - (contact hmodzelewski@eos.ubc.ca for access)

Past software releases
When:

- 1 release in 2006
- 2 releases in 2007
- 1 release in 2008
- 1 release in 2010

Past software releases: Categories:

- Wavefield reconstruction
- Wavefield separation
- Modelling
- Imaging
- Solvers (ISTc, SPGL1)
- Transforms (Curevelets, Surfacelets, ...)

- see Web Pages for more details

SLIM

Next software release

- We plan the next release in early 2011
- We intend to unify the software release as much as possible to:
  - clearly separate algorithms form applications
  - package together as many external software packages as possible
  - include automatic dependencies/unit tests
  - streamline installation and implementation of our software in the companies

SLIM

Next software release

- We plan the next release in early 2011
- The software release will contain the following applications:
  - L1 migration with active contribution from surface-related multiples
  - Compress Sensing of Full-waveform inversion gradients
  - Curvelet Matched Estimation of Primaries via Sparse Inversion
  - Estimation of surface-related primaries with L1 inversion and informed blind deconvolution techniques (update)
  - ...

Conclusions

- We have:
  - built a large software framework for seismic-imaging applications.
  - implemented new algorithms in this framework.
  - build a number of low-level tools to support those algorithms.
  - successfully started efforts to introduce MATLAB as our main research and development environment.

SLIM

Conclusions

- We made a huge effort in streamlining/ accelerating algorithm development and trying to allow researchers to focus on research rather than software development, and we wish to pursue this goal.
- MATLAB offers an efficient alternative so do we are compelled to exploit it and use to our benefit.

Acknowledgements

SINBAD

NSERC CRSNG

This work was in part financially supported by the Natural Sciences and Engineering Research Council of Canada Discovery Grant (22R81254) and the Collaborative Research and Development Grant DNOISE II (375142-08). This research was carried out as part of the SINBAD II project with support from the following organizations: BG Group, BP, Chevron, ConocoPhillips, Petrobras, Total SA, and WesternGeco.