IWAVE: a Framework for Regular-Grid Finite Difference Modeling

William W. Symes

The Rice Inversion Project

Department of Computational and Applied Mathematics
Rice University, Houston, TX

symes@caam.rice.edu

TRIP 20th Annual Review, 30 March 2012

IWAVE is…

- A framework for finite difference simulation
  - Common services – memory, communication, i/o, job control
  - Prescribed interfaces – problem description, numerical schemes
- Applications written to the framework interfaces
  - Staggered grid acoustics with PML
  - Staggered grid isotropic elasticity with PML
  - ....
- Portable – C, MPI 1, OpenMP
- Modeling engine for migration, inversion
- Open source…

IWAVE Project Goal:

Enable construction of new modeling code – physics, methods, hardware & software environments – without starting from scratch

Outline

- IWAVE project
  - Origin in SEAM QC
- Overview of the framework
  - Design principles
  - Grid data structures and automated exchange
  - Model definition interface
  - Job control
  - Performance & Ports

SEAM Phase I Acoustic Modeling

- Vendor: Tierra Geophysical (now Halliburton)
- 200 TB delivered 2010 – 60K shots, 500K traces per shot
- For the details: new SEG e-book by Fehler & Keliher
- Many insights into project management, QC – role of IWAVE

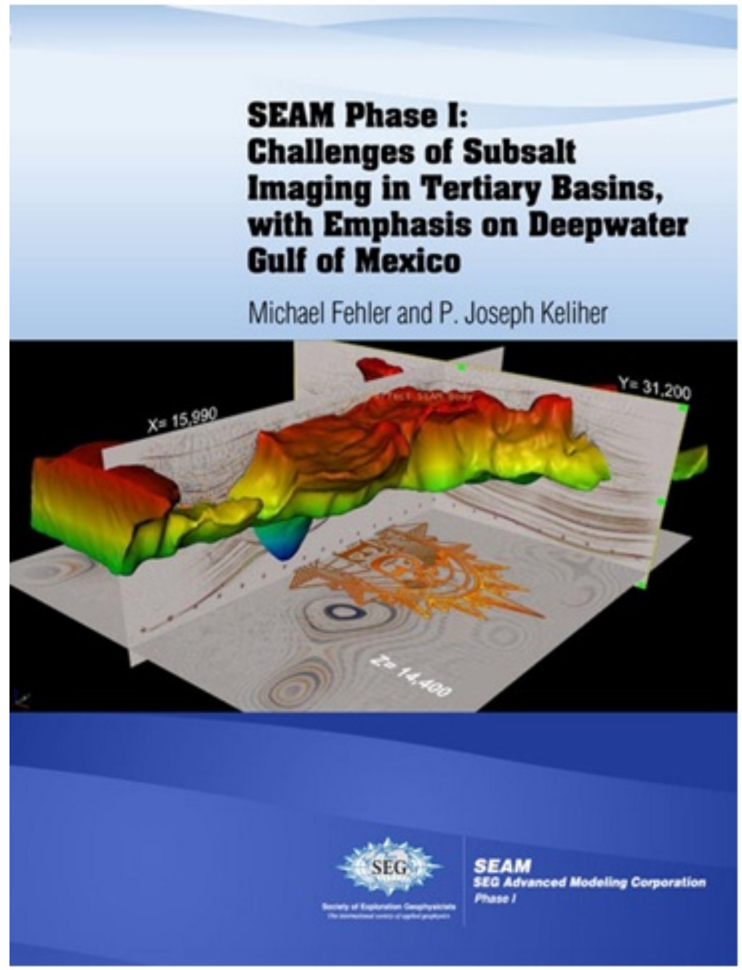

QC by spot-check

QC concept: use independent, open source, verified modeling code to spot-check vendor results

Requirements for SEAM Phase I GoM modeling task:

- 500K traces per shot, 16 s, 30 Hz, 8 ms, 2K samples
- 28x28x15 km modeling domain, typical sediment, salt properties ➔10 m grid, 50 GB per field, 40K time steps
- Calibrated direct wave field, specified pulse at calibration trace

Implications for simulator:

- Accurate 3D acoustic modeling with minimal gridding requirements (high order FD/FE)
- Effective absorbing, free surface BCs
- Parallelization via domain decomposition
- Scalability to 1000's of processes

The Origin of IWAVE – SEAM QC

Search by SEAM numerics committee: no available public domain code met specs

TRIP proposed to supply one…

- IWAVE 1.0: SEG 2009 (current 1.5)
- SEAM QC shots computed 2009-10
  - Ranger (Sun Opteron cluster, 60K cores) Texas Advanced Computing Center, University of Texas-Austin(2,10) acoustic staggered grid FD; 1-4K cores, subdomains (usually 2K); 6-12 hrs wallclock

Visual Comparison: Shot 20433, E-W line at N 10.585 km

Tierra

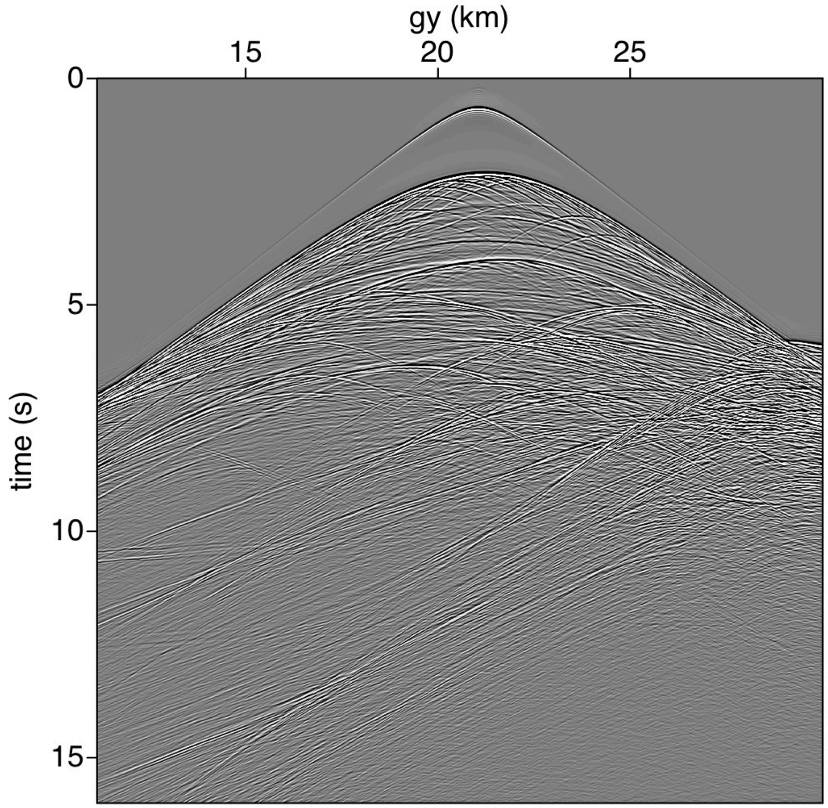

Visual Comparison: Shot 20433, E-W line at N 10.585 km
IWAVE

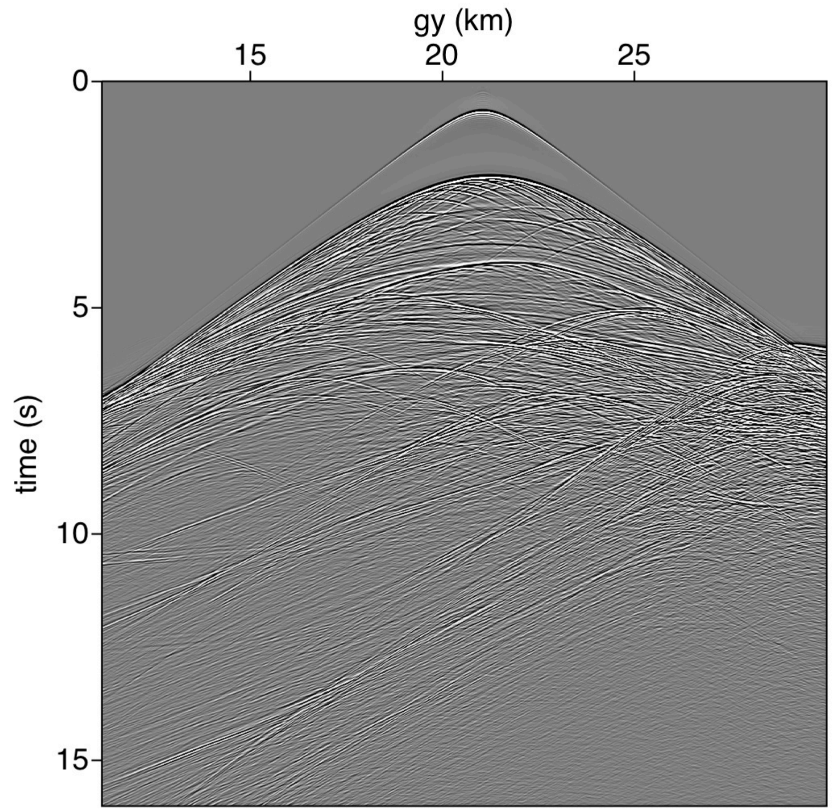

Why a framework?

- FD, FE apps share many common tasks:
  - Grid allocation
  - Data exchange patterns (domain decomposition) – depend on scheme
  - i/o w. common file structures (SEGY/SU, RSF,…)
- Many of these reusable across many apps, given interfaces and task def’ns
- Unstructured-mesh CFD- and FE-oriented frameworks: DUNE, deal.II, FEniCS, PETSc, Trilinos,…
- Regular grid FD/FE TD: restricted domain - additional opportunities for re-use, efficient implementation

IWAVE: Designers

- Igor Terentyev: overall design, core & utility modules, parallelism, staggered grid acoustic modeling (CAAM MA 2009)
- Tanya Vdovina: acoustic source calibration, PML for acoustics, verification
- Xin Wang: extensibility, staggered grid isotropic elastic modeling + PML
- Dong Sun, Marco Enriquez: RTM/FWI extension ("IWAVE++")
- WWS: various…

Support: SEAM Project, NSF, Sponsors of TRIP

IWAVE: Design Principles

- ISO C99
- "Object oriented C": built around small set of C structs containing data and functions operating on them
- Core modules: services – memory, comm, i/o, job control
- Application modules ("apps") – choices of physics, scheme
- Parallelism via domain decomposition, MPI
- Other forms of parallelism – initially OpenMP, task level via MPI, …
- Reads, writes data in standard exchange formats: initially SEGY/SU, RSF
- Open source, X11 license

IWAVE: Data Structures

RARR: rectangular subset of infinite lattice, plus virtual subarray info

•Allocated memory (real)
•Dimensional info: index tuples of axis origins in lattice, lengths
•Virtual subarray (pointer)
•Axis origin, length tuples for virtual subarray

Consequences:

•Every virtual array carries reference to its parent allocated array – can reshape
•Every virtual array located in common lattice – can compute overlaps, diffs

[IST, TRIP 2008]

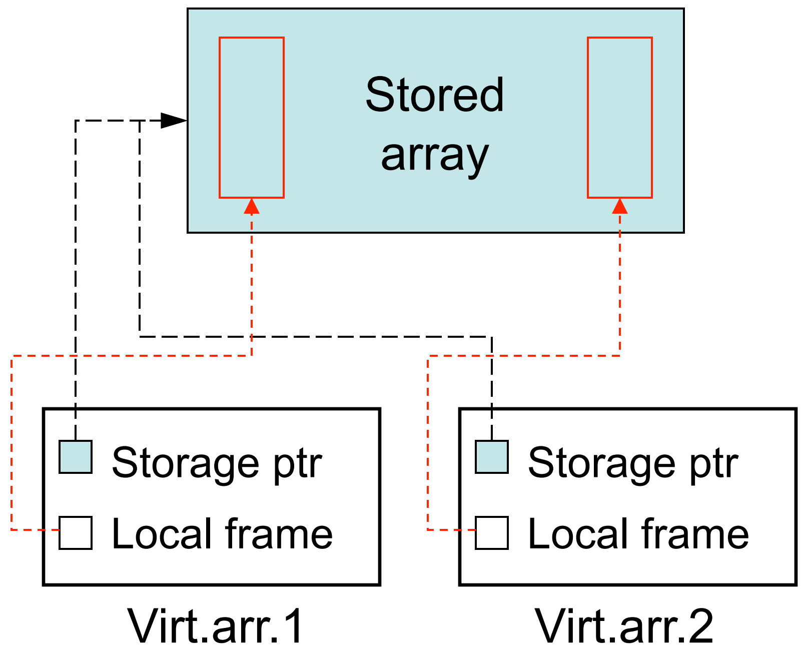

Methods:

- Create/resize subarrays
- Output
- Etc.

IWAVE: Data Structures

RDOM - collection of RARRs storing data for a model, scheme

FD_MODEL – functions to create and manipulate RDOMs, including time step

IMODEL = base (“allocated”) RDOM + RDOMS containing virtual computational, send, receive domain RARRs + FD_MODEL

Add driver to create modeling app

Defining a new IWAVE app

New FD_MODEL = functions with defined interfaces to:

- Identify dynamic arrays
- Define grid types – offsets from ref lattice
  - Expl: staggered grid schemes – vx(x+dx/2,y,z)
- Define stencil – dependency matrix, stencil masks
- Define sequence of updates, which fields updated in each substep in parallel
  - Expl: acoustics - update pressure, exchange pressure data, update velocities, exchange velocity data…
- Initialize opaque data structure needed in time step
- Time step function:
  - `ts_fun(RDOM * d, int iarr, void * fdpars)`

Job Control

Via parameter table, specifying info needed for app (eg. 3D sg acoustics):

```
-------------------------------------------------------------------------
FD info:
             cfl = 0.2
          cmin = 1.0
         cmax = 5.5
         fpeak = 0.015        central frequency
-------------------------------------------------------------------------
Model info:
     velocity = /work/00677/tg458297/SEAMDATA/VP.rsf
     density = /work/00677/tg458297/SEAMDATA/DN.rsf
-------------------------------------------------------------------------
Source info:
       srctype = point
        source = /work/00677/tg458297/SEAMDATA/wavelet.su
…
```

Job Control

Specifying domain decomposition, task parallelism:

```
-----------------------------------------------------------------
MPI info:
    mpi_np1 = 16        n_doms along axis 1
    mpi_np2 = 16        n_doms along axis 2
    mpi_np3 = 8         n_doms along axis 3
    partask = 1
```

Domains (approximately) uniform – include PML buffer zones

partask = number of shots to simulate simultaneously

Performance & Ports

Analysis at several levels for 2D, 3D acoustic staggered grid app

- Parallel – strong, weak scaling
  - see I. Terentyev 2009 TRIP AR
- Node/Multicore
  - talk by Adhianto & Mellor-Crummey this PM
  - new architectures, also SIMD vector instruction sets

Multicore performance

Platform: 192 x 12 Intel Westmere 2.83 GHz, 48 GB/node, QDR Infiniband, icc

Task: 12 shots, processed simultaneously, 2D acoustic sg (2,4) scheme, Marmousi model (751 x 2301, 5000 time steps) [D. Sun]

Upshot:
(1) at this level of scaling, parallel speedup is 100%, even for DD;
(2) memory contention (bus capacity? cache misses?) results in 3x slowdown on fully populated sockets, even when DD comm goes off-node

Nodes	PPN	Decomp	kernel
1	12	1 x 1	401 s
12	1	1 x 1	138 s
12	6	2 x 3	67.4 s
72	1	2 x 3	18.7 s

Early Experiences in Porting IWAVE to OpenCL on the Fusion APU

Ted Barragy, AMD

Bill Symes, Rice University

- HPC in Oil & Gas 2012
- March 1, 2012

AMD

The future is fusion

High Performance Compute Platform Based on multi-core DSP for Seismic Modeling and Imaging

Presenter: Murtaza Ali, Texas Instruments

Contributors:

Murtaza Ali, Eric Stotzer, Xiaohui Li, Texas Instruments

William Symes, Jan Odegard, Rice University

Thanks to

- IWAVE team – Igor Terentyev, Tanya Vdovina, Xin Wang, Dong Sun, Marco Enriquez
- Max Deschantsreiter, John Anderson, Scott Morton, Christof Stork, Ted Baragy, Murtaza Ali, John Mellor-Crummey, John Washbourne
- SEAM Project, NSF, sponsors of TRIP

www.trip.caam.rice.edu/software