

“Video Games” People Play at Work

(and other excuses to buy the best graphics chipsets possible)



Ryan Schneider, CTO, Acceleware Inc.

Get superpowered!



www.acceleware.com

About...

- Acceleware
 - 3 years old, 30 people, growing!
 - Canadian ... Calgary (-5°C, 23°F), Tampa (25°C, 77°F)
 - Shipping NVIDIA GPU's for commercial computing for ~2 years
- My Background
 - Hardware acceleration of EM simulations using FPGA's
 - Digital Software Radio – DSP + FPGA + RF Hardware
 - Integer arithmetic, VHDL floating-point units, mixed precision
 - High Performance Computing
 - 87 node cluster, MPI, NUMA







Outline

- A New Era of “Games”: Target problems and markets
 - Context
 - Electromagnetics (Demo...)
 - “Energy” – Seismic Processing, Reservoir Simulation
 - Conjugate Gradient Solver
 - Performance Results
- Working with GPU’s
 - Multi-GPU Computing
 - Nvidia Quadroplex
 - Challenges of GPU’s
- The Future...



High Performance Computing Markets

- **Electromagnetics (EM)**
- **Energy**
- Biomedical
- Computational Fluid Dynamics
- Pharmaceuticals
- Industrial
- Military



High Performance Computing Markets

- *Electromagnetics (EM)*
- *Energy*
- Biomedical
- Computational Fluid Dynamics
- Pharmaceuticals
- Industrial
- Military

Common Theme:

1. Compute-Intensive Tasks
2. (On Traditional Hardware) Tasks Will Run for Hours, Days, Weeks

GPU's as SuperComputers?

Larger Context

- MANY hardware choices – Cell, Stream Processors, FPGA's, Ageia, Stretch, CPU in Memory, Clearspeed

GPU's are great for CERTAIN problems

- (Not) Cryptography
- (Not) E-mail Acceleration

Acceleware Solutions: Three Prongs


1. Algorithmic Tuning/Enhancements
2. Large Memory Bandwidth
3. Very Parallel Computational Resources



Electromagnetic (EM) Simulation “Games”

“Game” Physics


- Finite Difference
 - 3D Data Sets (Volumes)
 - Nearest Neighbor Look Up
 - 32-bit Precision
 - More complicated materials, boundaries
- Finite Element




$$E_z^{n+1} = E_z^n + (H_x - H_x) + (H_y - H_y)$$



Electromagnetic (EM) Simulation “Games”



Next-gen Pacemaker



Next-gen Hearing Aid




“Graphics” (10 out of 10): Extremely *Realistic* Modeling of:

- Head-Cellphone Interaction
- Photonics
- Lithography Masks
- MRI development/modeling
- Printed Circuit Boards

“Game”play (0 out of 10): Without the right graphics chipset, one “level” take hours to days.



Electromagnetic (EM) Simulator Benchmarks



Either 1, 2 or 4 “boards”
connected in one
workstation



DEMO




www.acceleware.com


Acceleware Energy Products “Oil Derrick Tycoon” Game



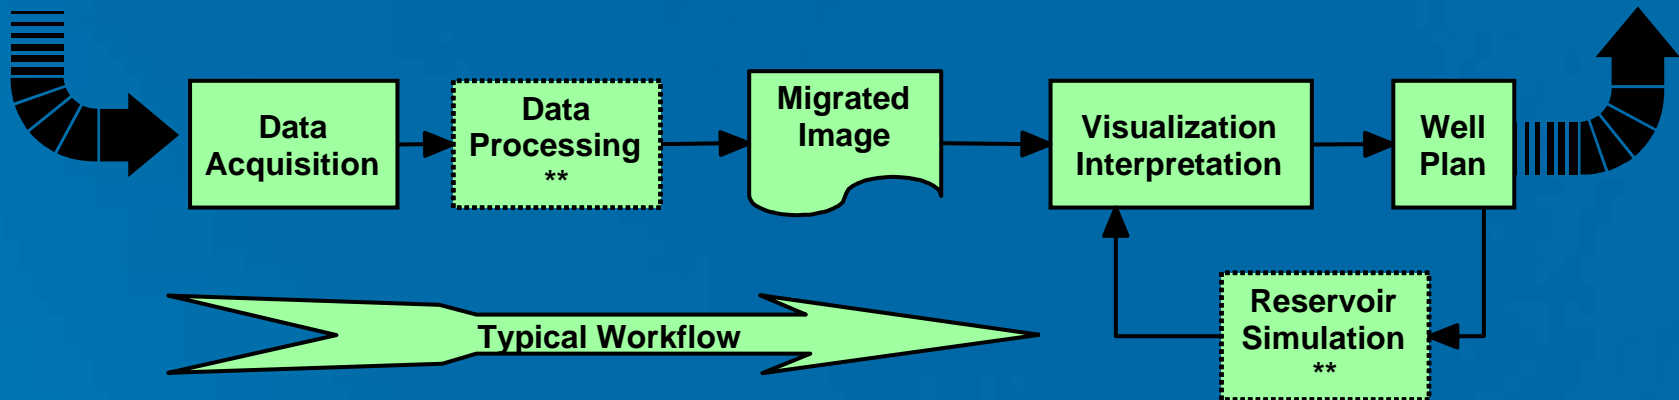
“Oil is found in the
Mind”



Confidence is in the
details




Confirmation is in
the drill bit



** Most compute-intensive
components




Acceleware's "Oil Derrick Tycoon" Seismic Data Processing



"For every 10' away and every 1' down
aim an additional 4" below the target"

Seismic data processing
is any
computational technique
that attempts to
remove noise
or
wave propagation effects
in order
to recreate
an image of the
subsurface.

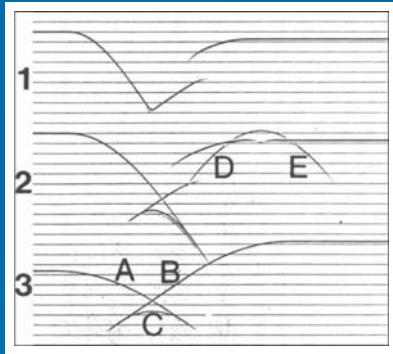


Earth Stratigraphy demands
more intense algorithms

Migration and pre-stack imaging are dominated by
Kirchhoff integral equation methods.




Acceleware's "Oil Derrick Tycoon" Seismic Data Processing




Synthetic model
before migration




Migration is the most computational intensive component. Migration moves dipping reflections to their true sub-surface positions and collapses diffractions



Synthetic model after migration



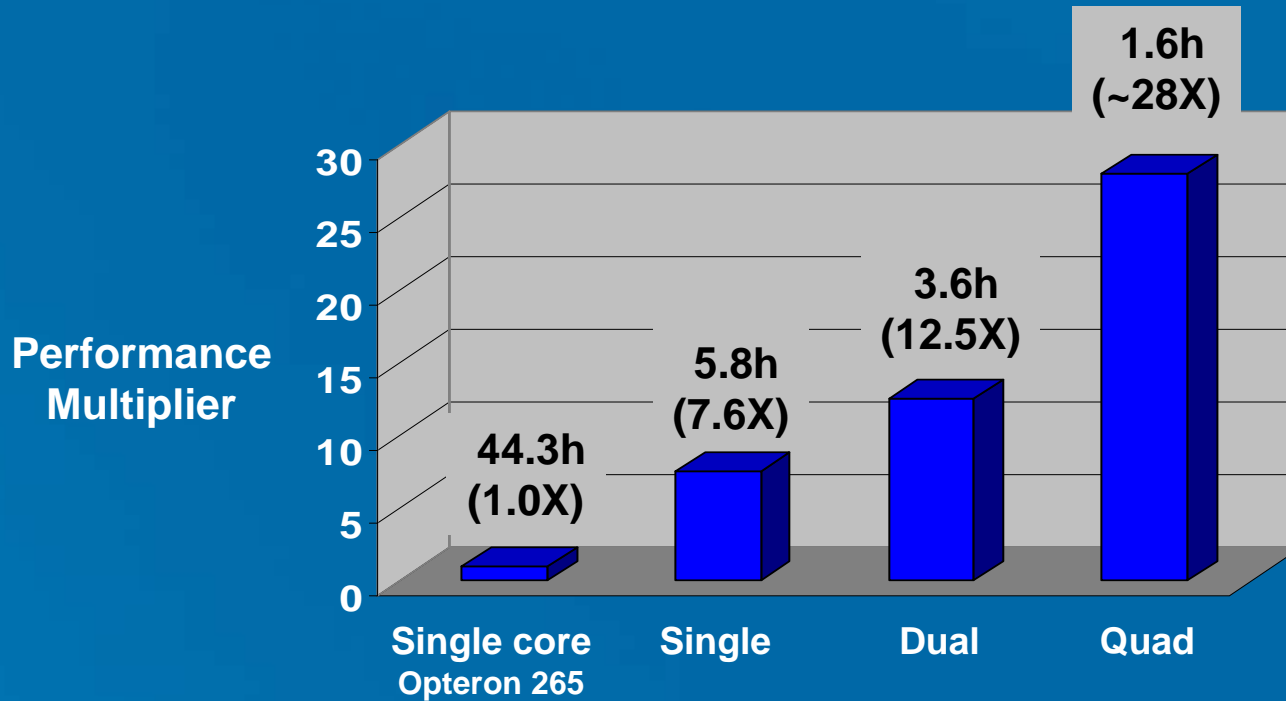
Seismic section before migration



Seismic section after migration

Single Seismic "Job": Weeks to (marine) 1.5 years.

Acceleware's "Oil Derrick Tycoon" Game Kirchhoff Migration Performance



Each additional GPU ~ 7X



Matrix Inversion “Games” Conjugate Gradient Solver


$$[A][x] = [b]$$

- Sparse
- Simplest conjugate gradient, no pre-conditioning
- 32-bit
- 3X over optimized Intel/AMD
- Future: Mixed-precision...



Multi-GPU Computing

- EM
 - Domain Decomposition
 - Manual Synchronization
 - Host as Intermediary
 - Double Transfers!
 - Linear scaling for 2 and 4 GPU's, 6.5X for 8 GPU's
- Linux -> X11 Screens



QuadroPlex

- Advantages

- Performance
 - Same as Dual-card Workstation
- Form factor
 - Own Power Supply
 - 1 PCI-E Slot Equals 2 PCI-E Slots
 - Compact, Rack-mountable
 - *Existing Installs*
 - “Peripheral”

- Disadvantages

- Price Not Competitive For “Compute”
- Platform Support



“Old” Challenges of GPU’s

Challenge	Solution
Programming Model	CUDA, Abstraction, Encapsulation
High Overhead	CUDA ...
Remote Access	Fixed ...
Memory Capacity	Multi-GPU, (Samsung?)
Lack of Documented Architecture	...



New Challenges of GPU's

- Acceleware: Technically capable of delivering GPU solutions into the marketplace...
- **CALL TO ACTION**
 - Drive Adoption – GPU's in the Corporate IT Model
 - High-End Graphics Chipsets in **Every** Workstation
 - Memory Capacity
 - Double Precision



Summary: More Gaming at Work!

- GPU's for Commercial Computing Now!
 - 5-10X speed up per GPU
 - Various applications
- Progress in Tools, Architectures
 - Discontinuous, order of magnitude increase in compute capabilities
 - New opportunities – This is the beginning!
- Acceleware end users: Out-computing their competition!



**Support Commercial Computing!
Buy a copy of Doom 3 today!**

