

Connections between seismic imaging and high-performance computing

John T. Etgen
Distinguished Advisor, BP

Agenda

Brief seismic imaging primer

Trends

- Geophysics
- Computing and how we interact with it

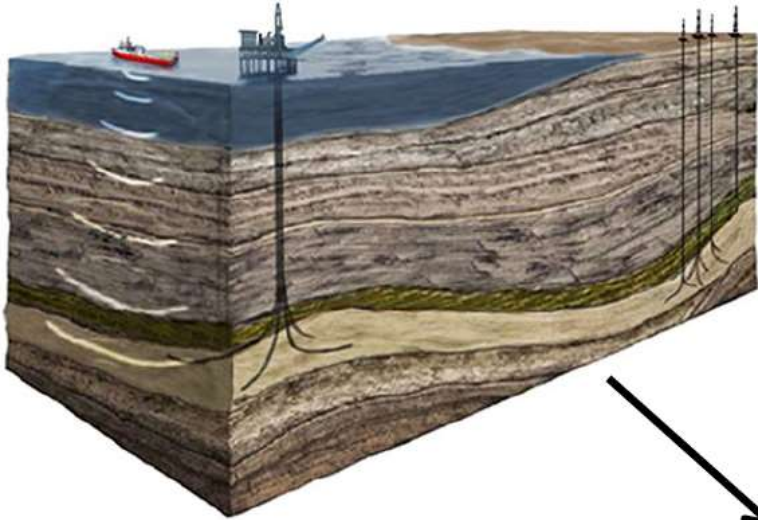
Gaps

- Graphics
- Skills
- The lack of an “industrial revolution”

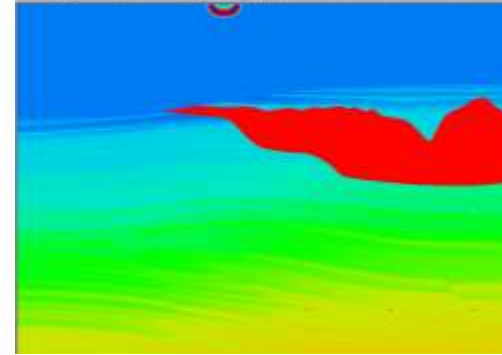
Some observations for the future

Industrial seismology: brief intro


We transmit sound waves into the Earth and record the echoes



Sedimentary rocks are layered and reflect sound waves from interfaces between layers



We process the recorded echoes to make images of the subsurface that look like this



Geophysics trends

Main imaging / processing algorithms staying relatively stable

Data density going up (even w/ downturn in the industry)

Continued desire for improved resolution drives computational effort non-linearly: $(\text{frequency})^4$

Algorithm trends

Acoustic wave propagation for RTM and FWI dominate

- ✓ Hyperbolic PDE: Mostly explicit time marching methods, some FFT-based, some convolutional
- ✓ Some Helmholtz aficionados still out there

Time domain elastic wave propagation growing

- ✓ Some FD, some FE, some SpecE

“Old stuff” still hanging in: Kirchhoff, tomography, data processing, sorting

Linear and quasi-linear optimization

Some uncommon tasks for us

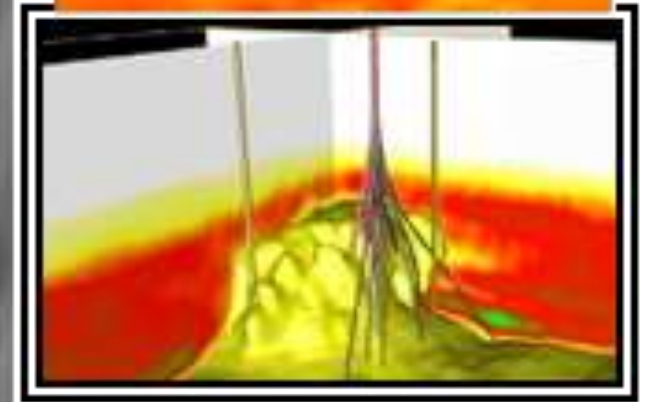

- × Exact solves of large linear systems...(frequency domain FD is the exception)
- × Searching large databases
- × Double precision arithmetic / heavy use of integers
- × Unstructured grid methods
- × Fully non-linear optimization

FWI increasingly contributing to imaging

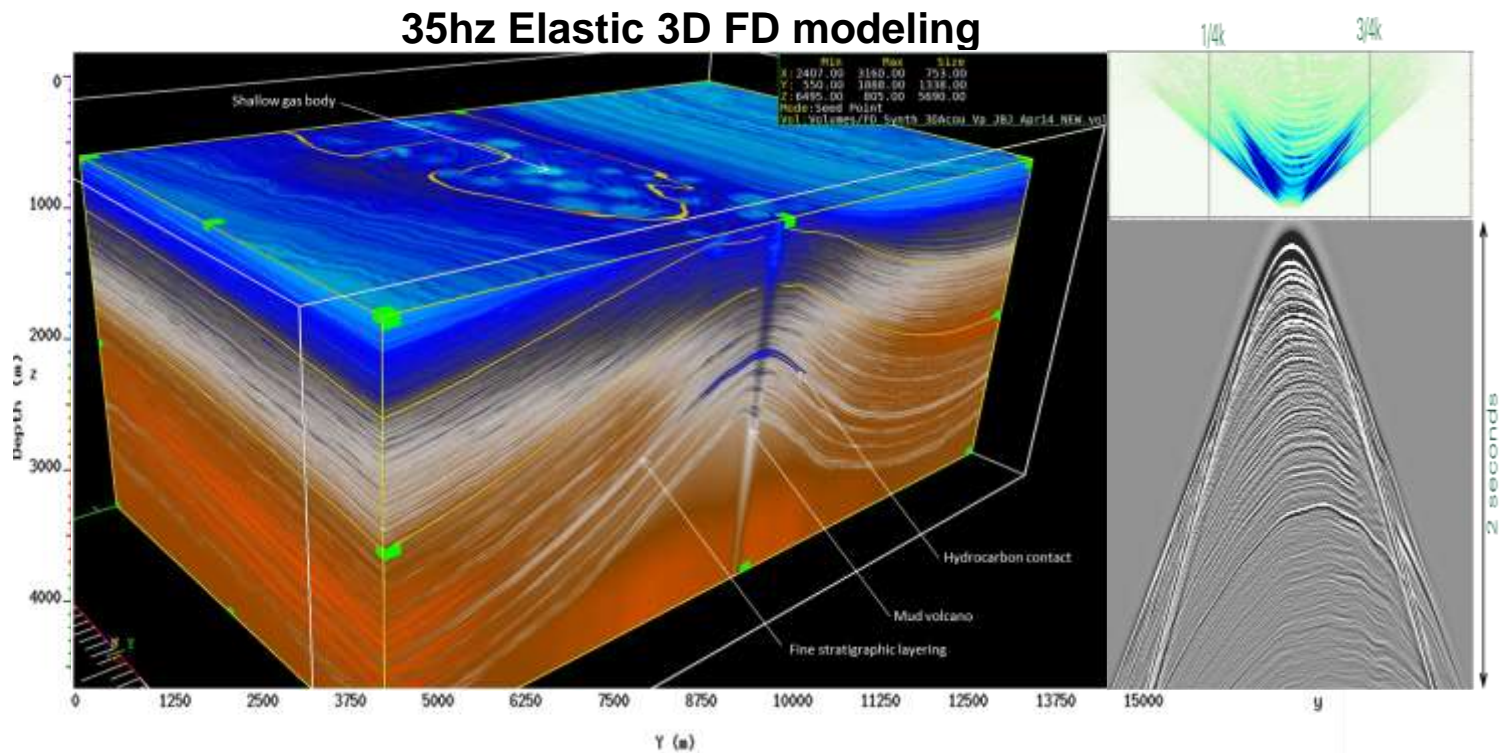
Conventional



Full Waveform Inversion

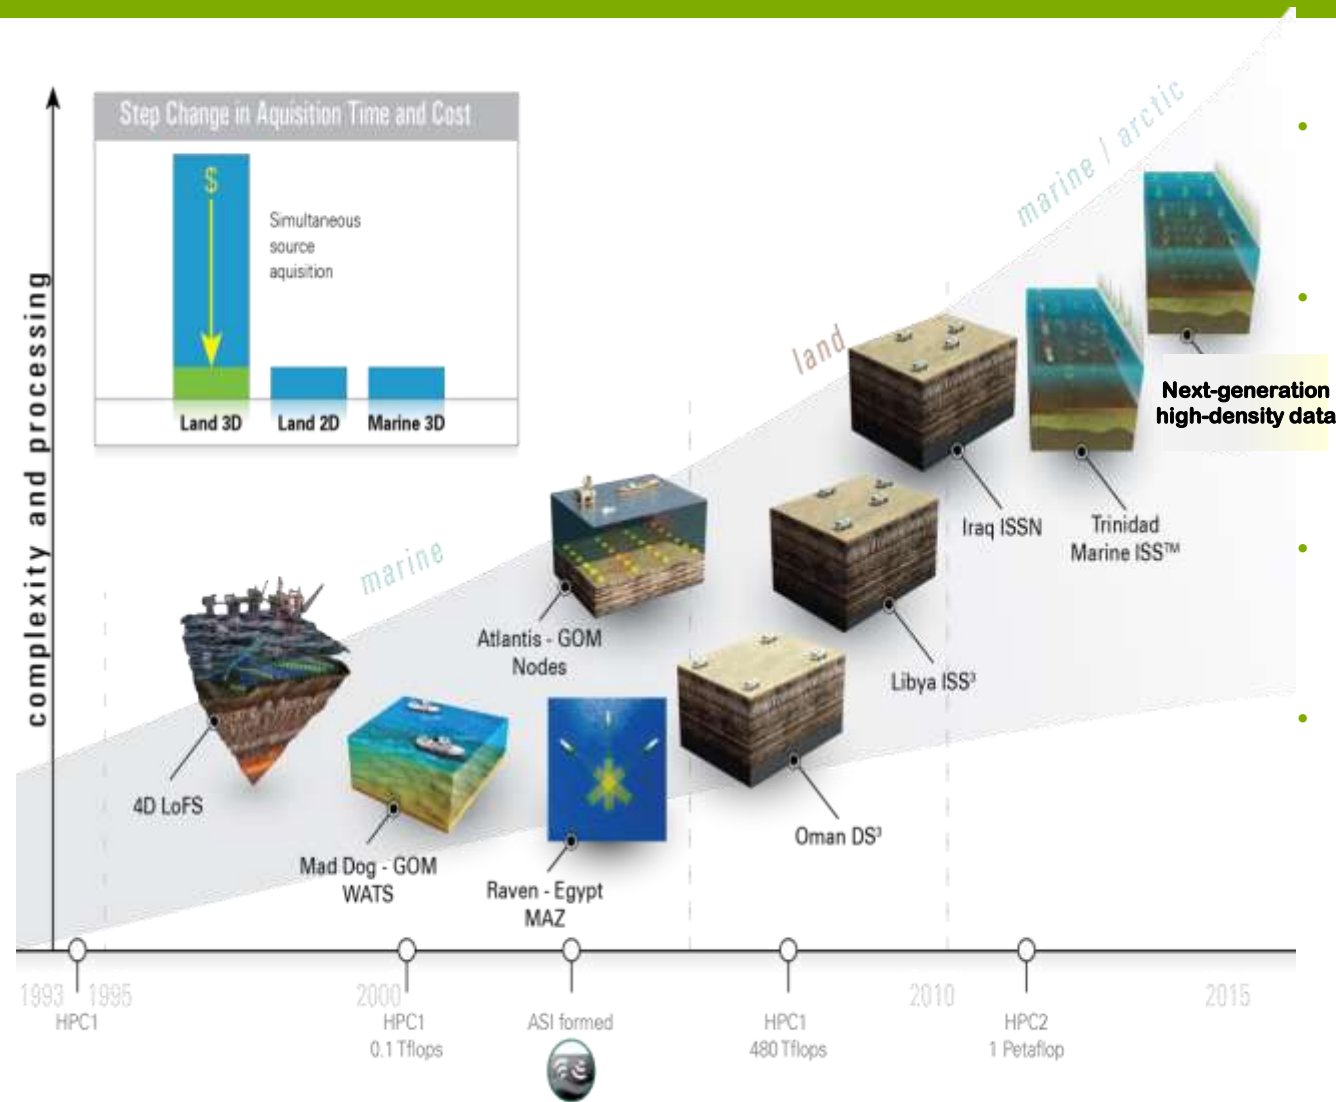


Large (industry) scale elastic modeling



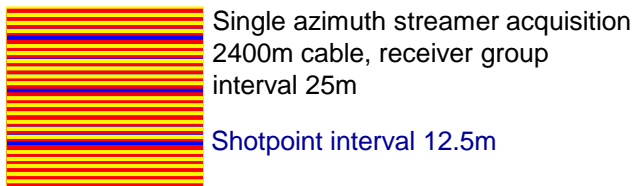
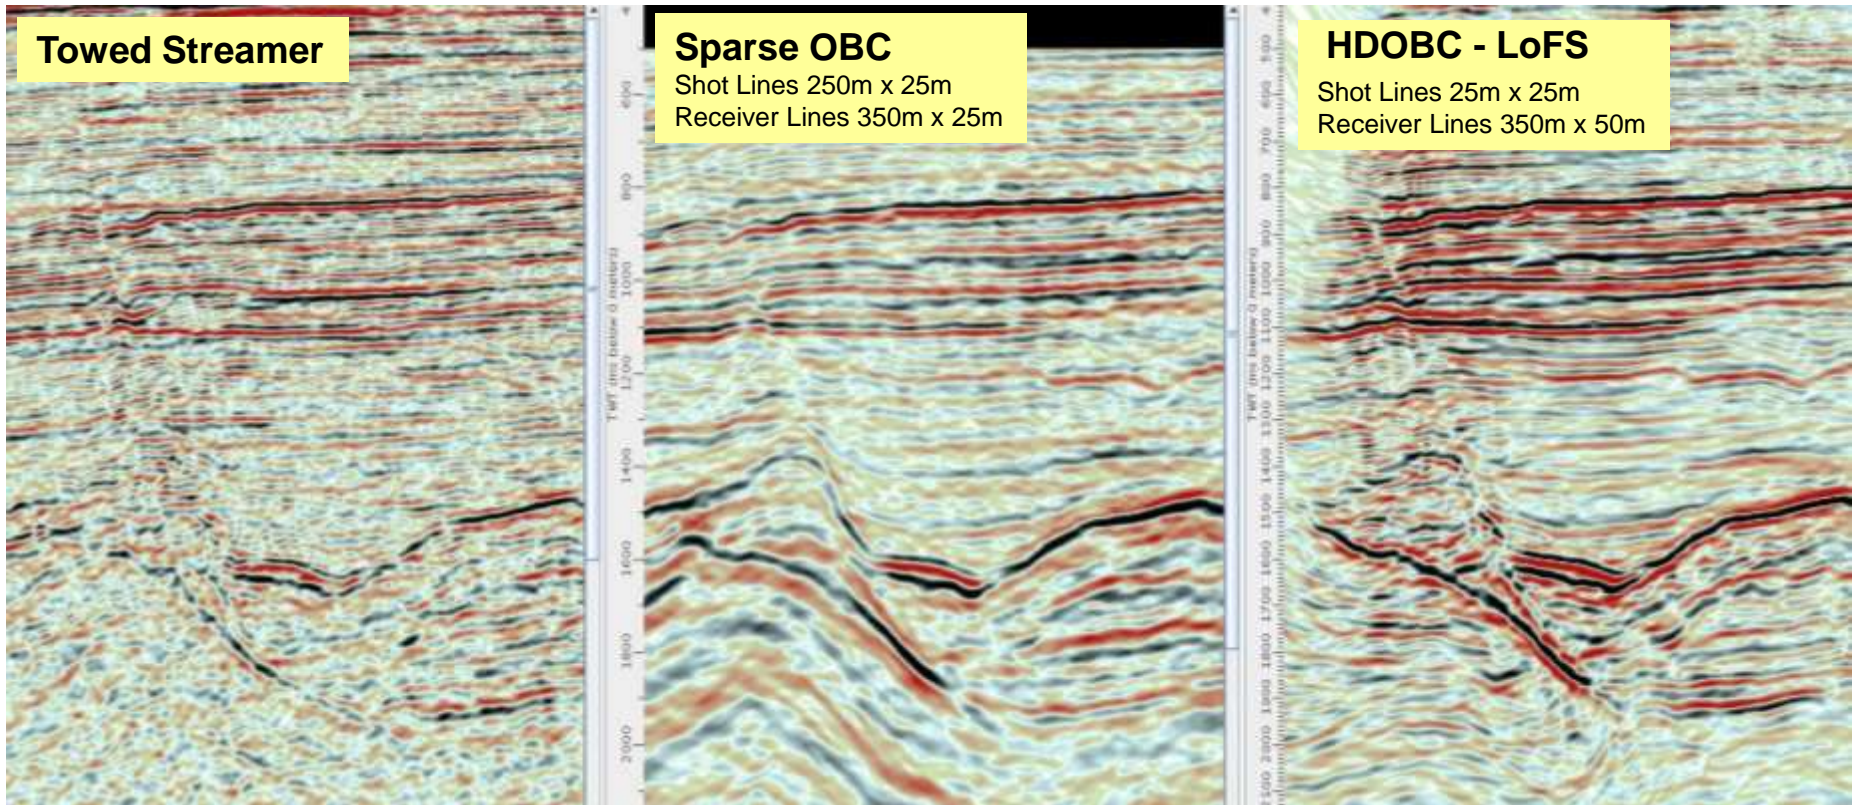
FD data generation: 86,928 teraflop-days

Data density increasing as sensing becomes cheaper and need for resolution and S/N becomes more urgent



- Marine data density 20 yrs ago: 160,000 traces/km²
- Marine data density now: >6,000,000 traces/km²
- Land data density 20 yrs ago: 43,000 traces/km²
- Land data density now:
Expl: ~1,000,000 traces/km²
Devel: ~40,000,000 traces/km²

High Density OBC (HDOBC) Data Quality



Trace density ~ 75K traces/Km²



Trace density ~ 400K traces/Km²



Trace density ~ 4M traces/Km²

after Davies et al, (2011) @Offshore Europe (conf.)

Computing trends and observations

I've seen at least 6 paradigms in 25 years:

minicomputers w/array processors, vector computers, fine-grained SIMD, RISC SMP, coarse-grained clusters, "the current state"

Geophysical algorithms must be trivial, as we make (made) effective use of all of these!

I have no clue what the next significant paradigm is

To communicate, you need to know why you invest in computing

Some of my personal biases

A good algorithm wins *every time*...

A fast enough computer beats a good algorithm *every time*...

More memory (total, per process, whatever) is always good...

Memory balance should be: 2 reads +1 write per Mult-Add!

A computer is an appliance, not a temple... use it as appropriate, don't worship it...

2 types of “HPC consumers”

Computer is part of my “brain”...

- Solve new problems, prove up new technology

- Code changes often (and that’s good)

- May be more willing to experiment with hardware, languages, paradigms

- Should not monitor “utilization” fanatically...

Computer is my money printer...

- Solve the “same problem” many times for \$\$\$

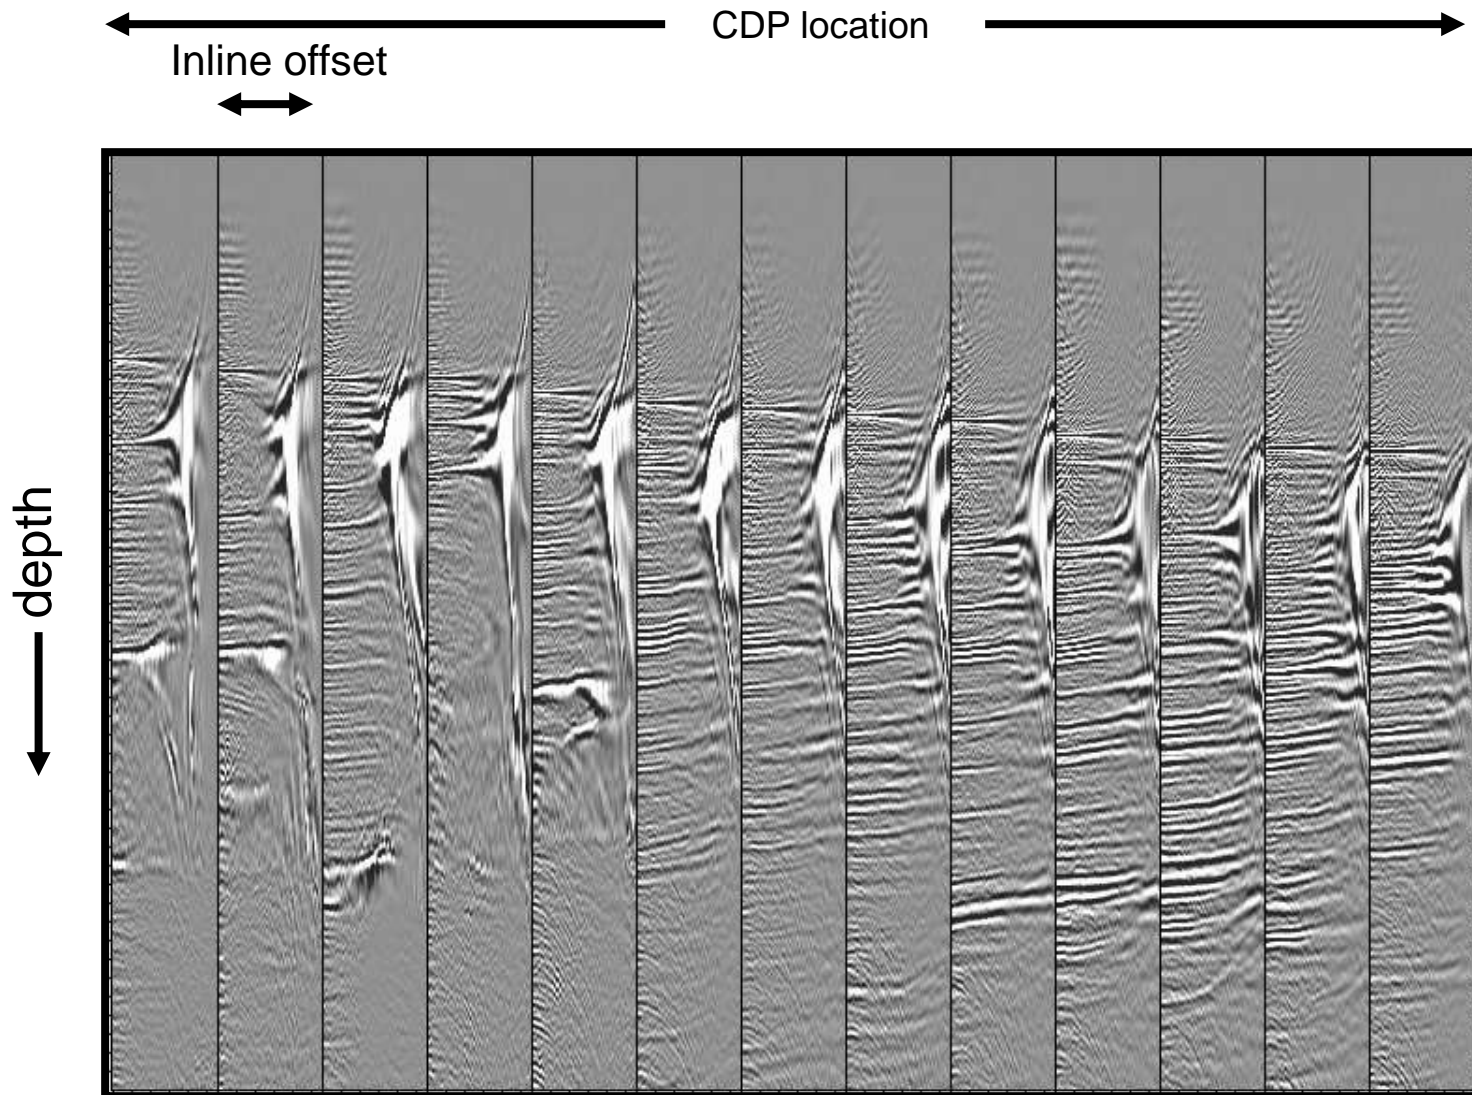
- Code changes more “managed”

- Willing to invest hard-coding time, but then your business is margin.

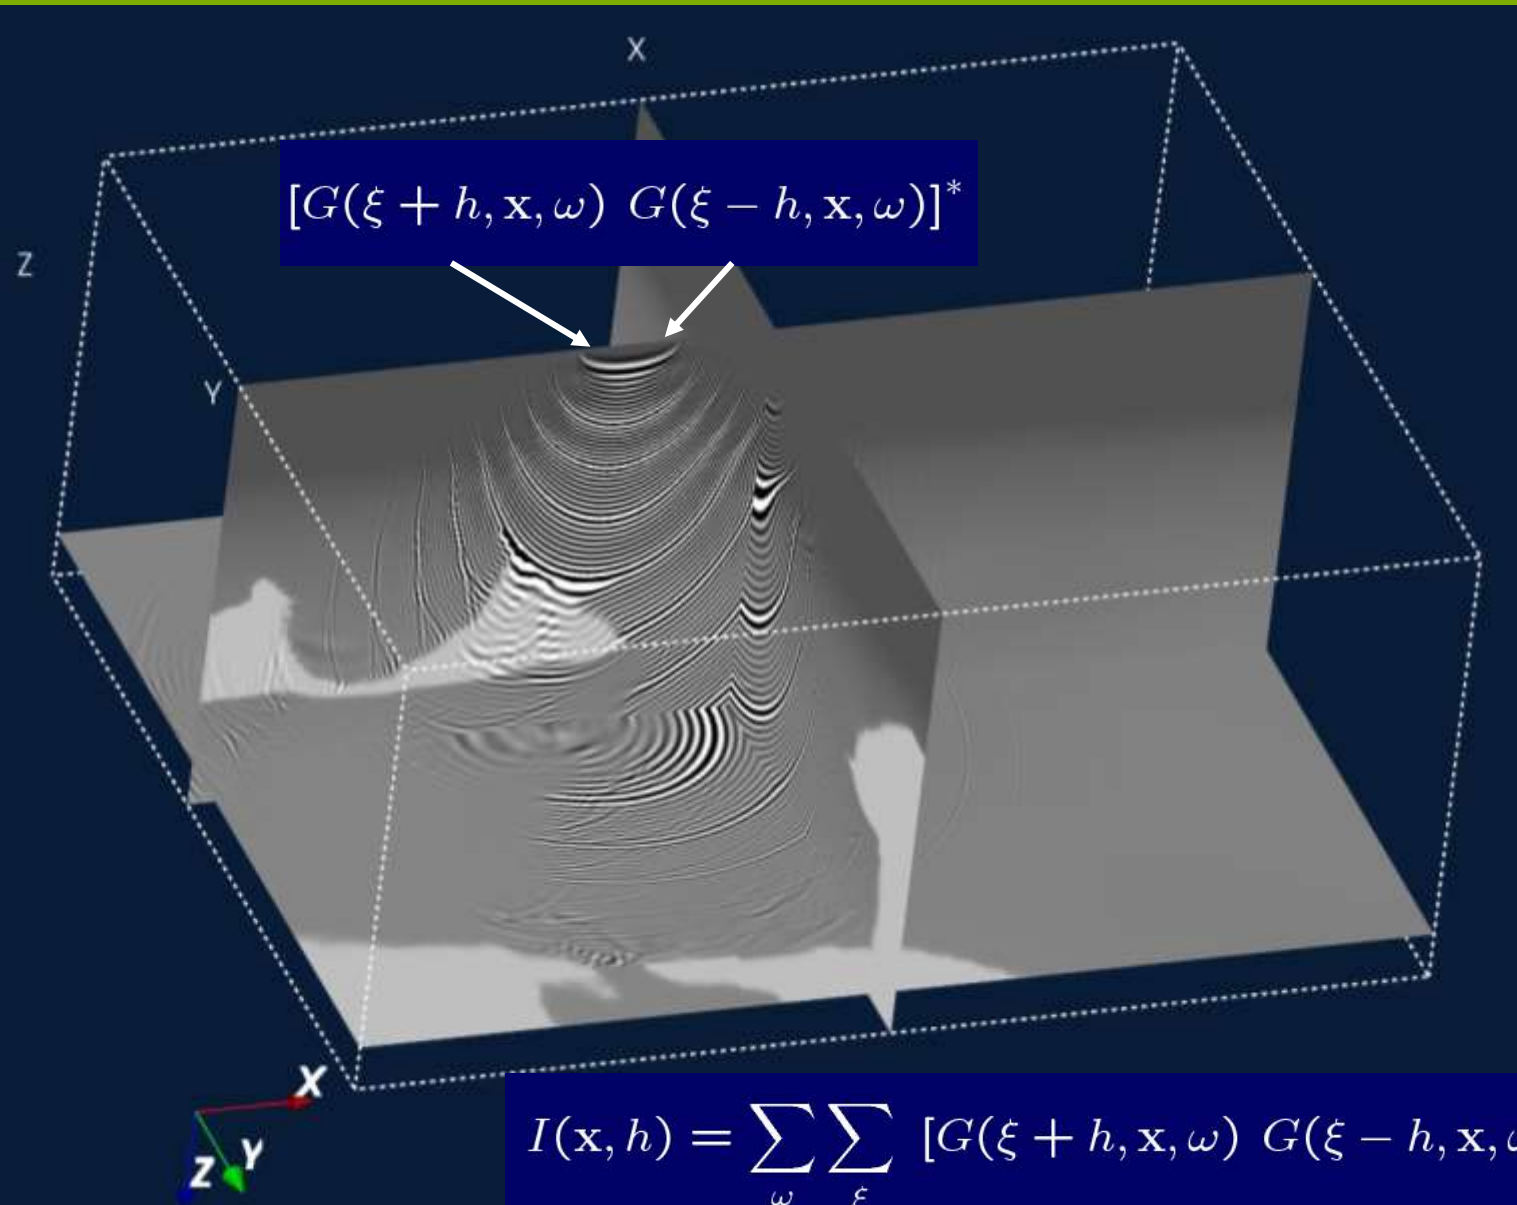
- Large fixed costs/initial costs easier to tolerate?

- You’ll hear about “utilization”

Graphics from the '60's



Graphics from... well... later anyway



Divergence of Skills

Note to self:

Rant about how hard it's getting to be good at applied science and computing at the same time...

And how hard it is to convince people to try...

And how hard it is to convince anyone that this is a problem...

And why do compiles take just as long now as they did 20 years ago?

“Industrial Revolution”

Note to self:

Rant about how we are not very good at building the total computational exercise out of interchangeable/standard “parts”

It always seems easier to “do it yourself”


Collaboration has not gotten any better, but it should have, as the complexity of our endeavor has exploded.

And what happened to static executables?

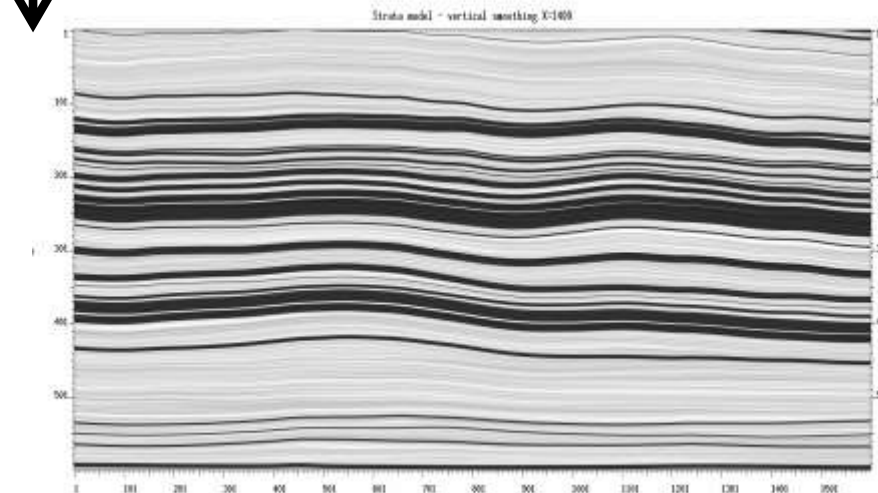
And demand-paged virtual memory?

Leveraging HPC for subsurface understanding


Geological understanding and analogs



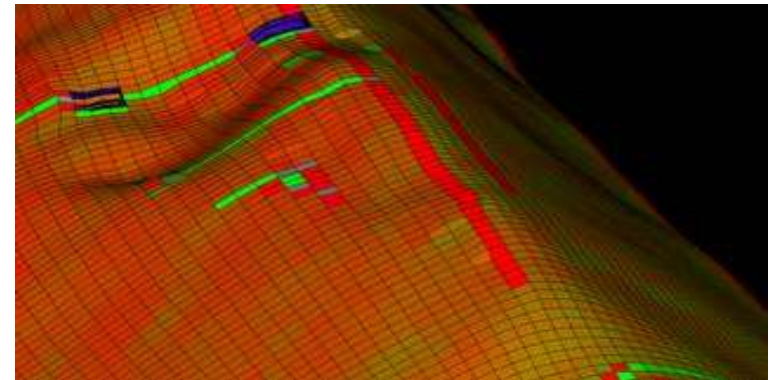
Statistical geology simulation



Elastic wave simulation



Reservoir simulation/ history match



What's really important

“Performance is everything” - F.H.

“If I could have one wish, it would be for a few more tools to help” -F.H.