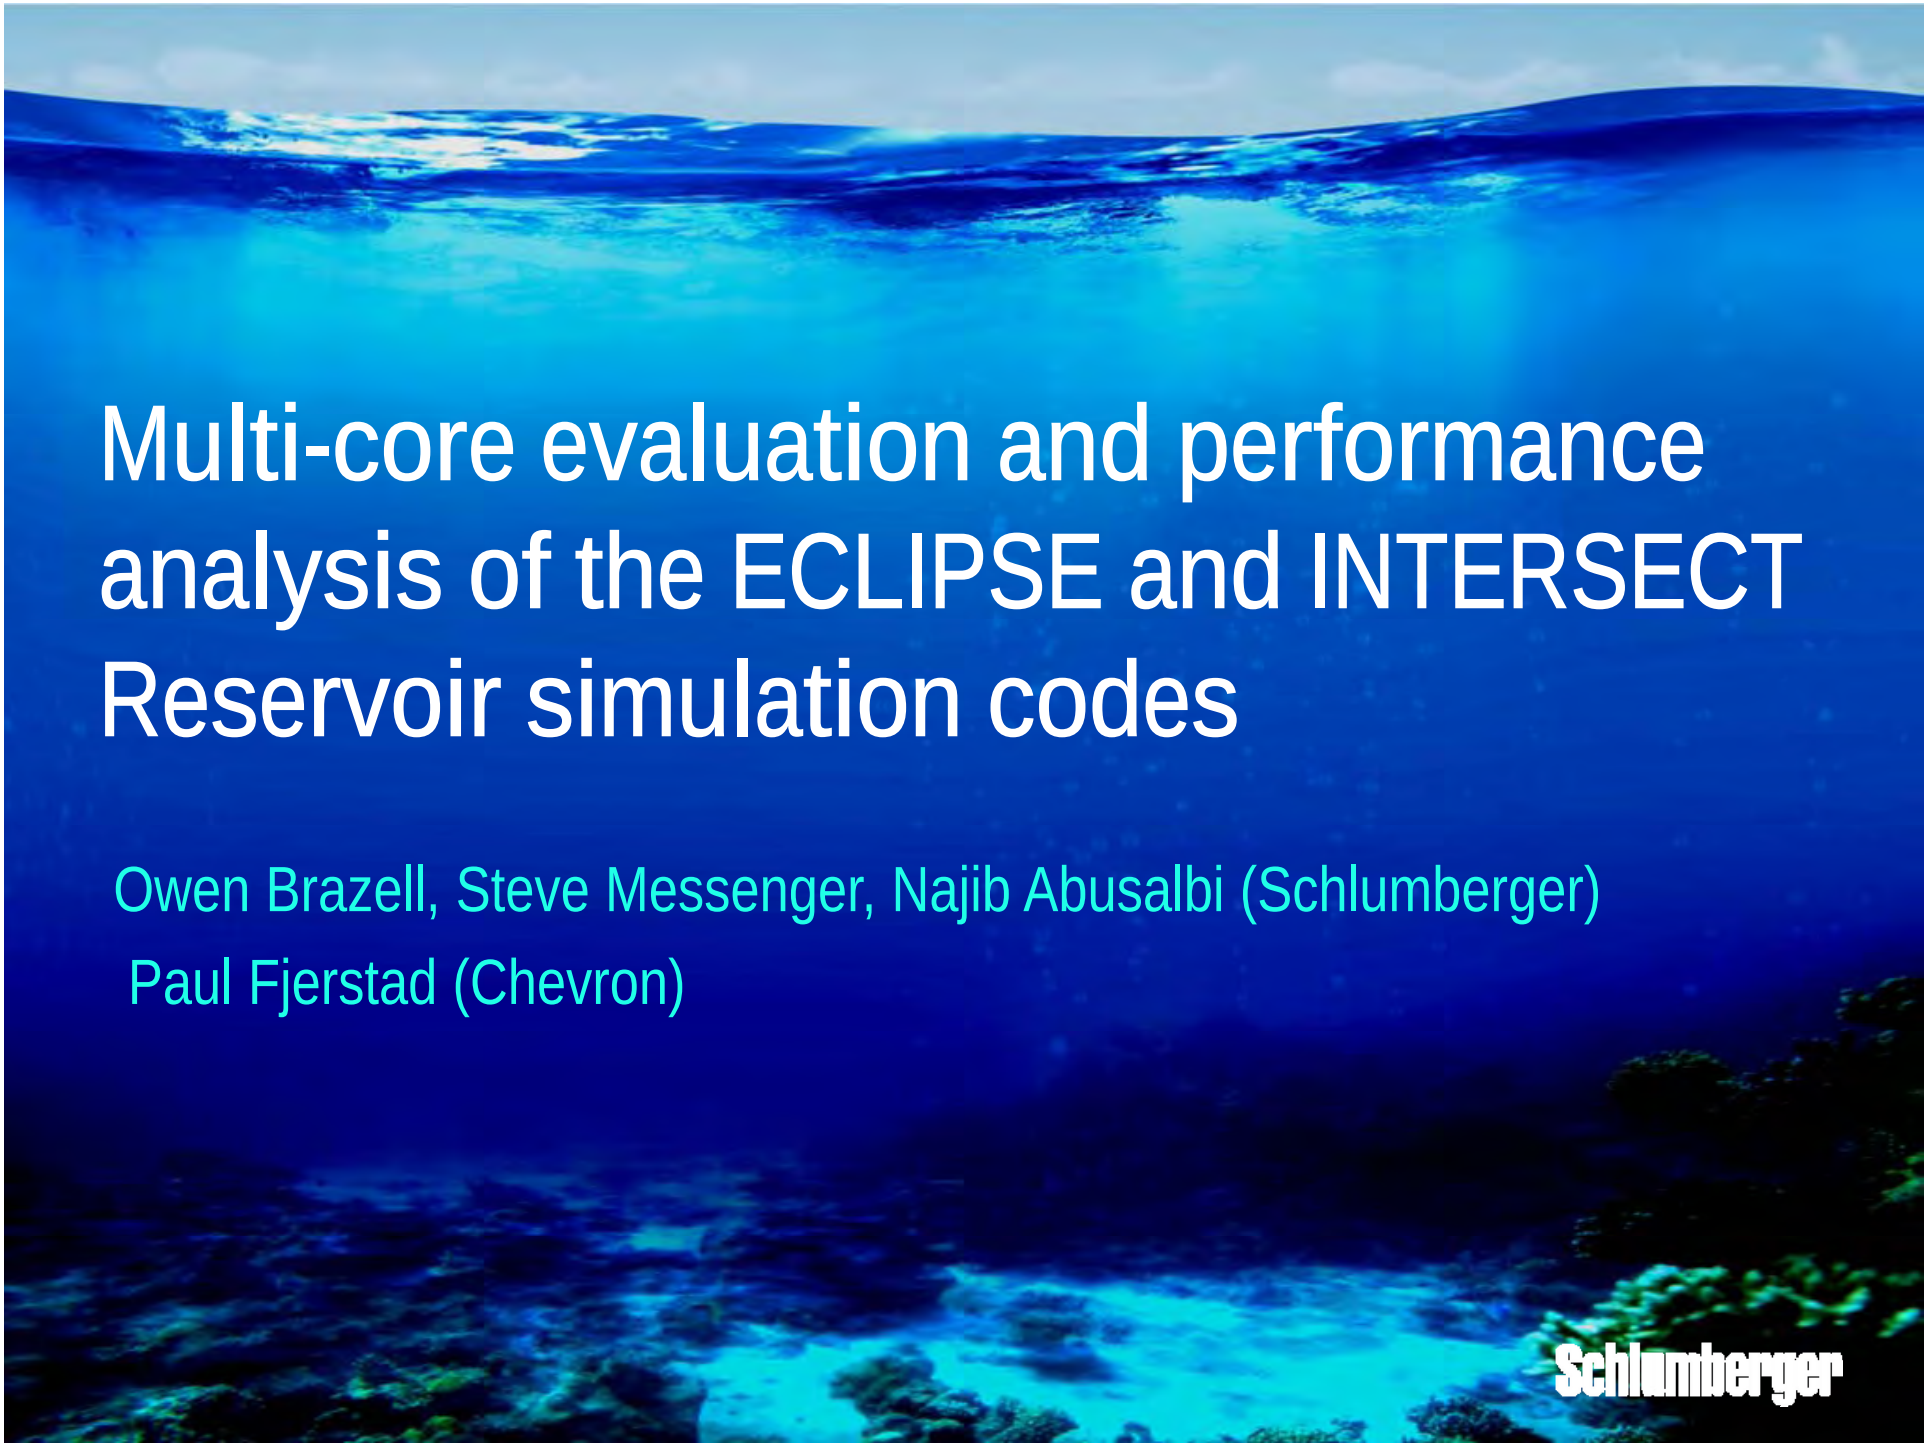
Multi-core evaluation and performance analysis of the ECLIPSE and INTERSECT Reservoir simulation codes
Owen Brazell, Steve Messenger, Najib Abusalbi (Schlumberger)
Paul Fjerstad (Chevron)
Schlumberger

Intel Legal Disclaimer

- All products, computer systems, dates, and figures specified are preliminary based on current expectations, and are subject to change without notice.
- Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing.
- Intel does not control or audit the design or implementation of third party benchmarks or Web sites referenced in this document. Intel encourages all of its customers to visit the referenced Web sites or others where similar performance benchmarks are reported and confirm whether the referenced benchmarks are accurate and reflect performance of systems available for purchase.
- Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. See www.intel.com/products/processor_number for details.
- Intel, processors, chipsets, and desktop boards may contain design defects or errors known as errata, which may cause the product to deviate from published specifications. Current characterized errata are available on request.
- Intel, Intel Xeon, Intel Core microarchitecture, and the Intel logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
-

Defining the challenge

- Need to understand future evolution of hardware
  - More cores per socket
  - Server cooling costs suggest that we probably need to keep the thermal envelope per socket the same therefore clock speeds per core will drop.
- Can the current generation of simulators take advantage of this?

Schlumberger SIS simulator suite

- Current generation of ECLIPSE simulator family
  - ECLIPSE 100 (black oil model, distributed memory model using MPI)
  - ECLIPSE 300 (compositional/Thermal model, distributed memory model using MPI)
  - FrontSim (Streamline simulator using Multi threaded options – limited to running on a single SMP system)
- INTERSECT (Next generation simulator , distributed memory model using MPI)

Need to understand performance of simulators in todays multi-core environment

Run benchmarks of the simulator suites on current and upcoming chips

- Intel® Xeon® Procesor 5500 series (4 cores,8 threads/socket)
- Future Intel Xeon processor, codenamed Westmere EP (up to 6c/12t)
- Future Intel Xeon processor, codenamed Nehalem EX (up to 8c/16t)
- AMD Istanbul (6 cores/socket)
- AMD Shangai (4 cores/per socket)
- AMD Magny Cours (12 cores per socket, not yet available for test)
- IBM Power 7 – not yet working well with ECLIPSE, working with IBM

Various simulation models used to check performance

- All models are the standard benchmark models provided with the ECLIPSE DVD
- Data sets that we use for our hardware certification program
- Work closely with Intel and others to make sure hardware is best tuned for the simulations

ECLIPSE 100 Blackoil model

- One million cell Black oil problem (artificial)

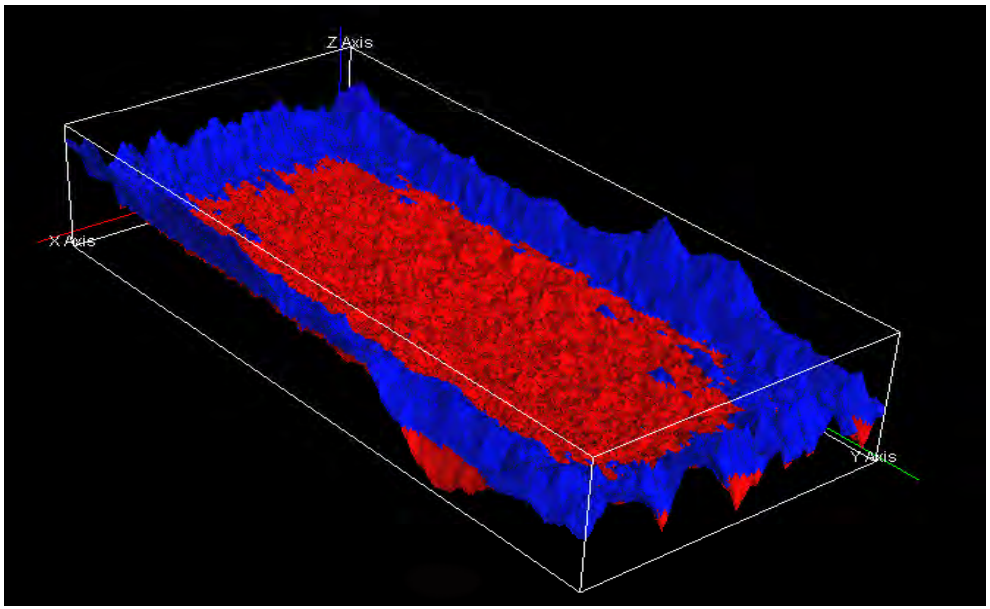

ECLIPSE 300

- 2 Million cells Compositional model (7 component) – based on a North sea field

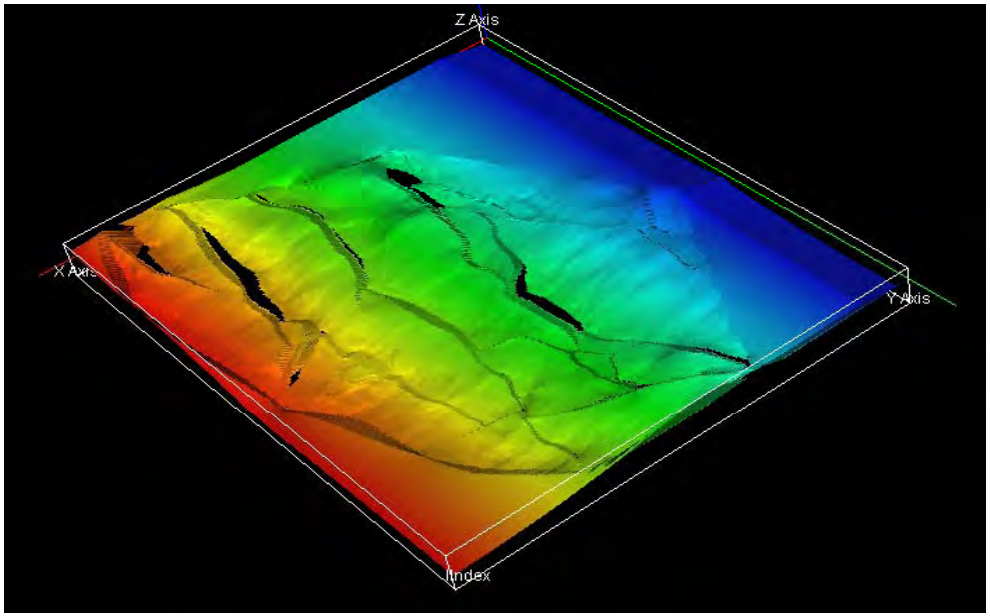

FrontSim – 448800 cells

- Synthetic model – variant on SPE10 model
- Designed to spend most time in Saturation solver to show scalability
- Pressure solver not multi threaded

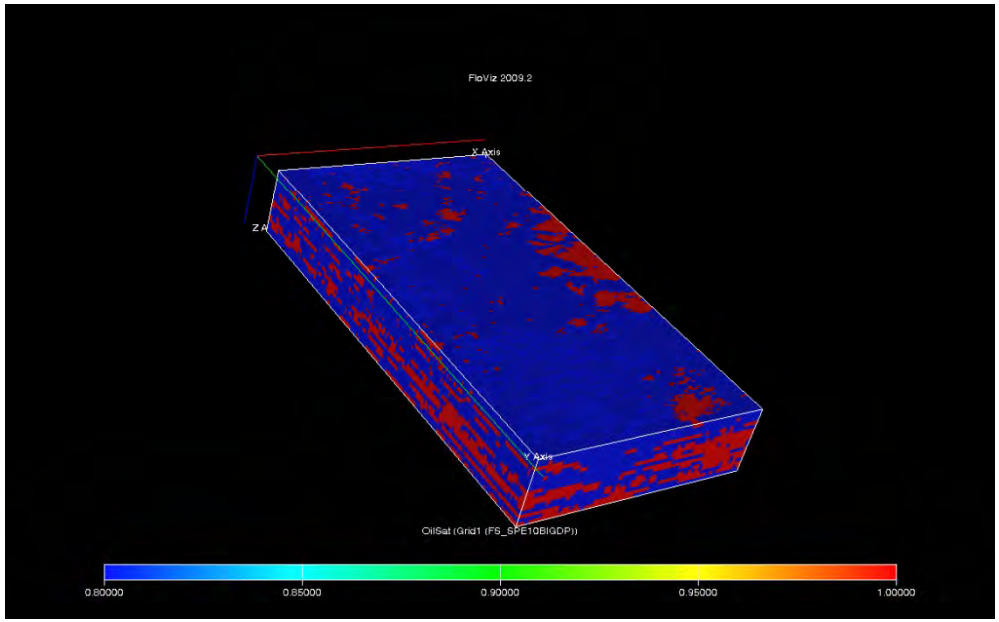

Hardware used in Benchmarking exercise

- Hardware used was early access hardware
- This hardware has not yet been optimized for final release
- Part of the reason for testing was the close relationship between Intel and Schlumberger and by doing these tests any issues found can be sorted out before release

Legal Information

Due to compliance with Intel pre-launch embargo rules we have blanked out the absolute timing data on all the slides

ECLIPSE scaling

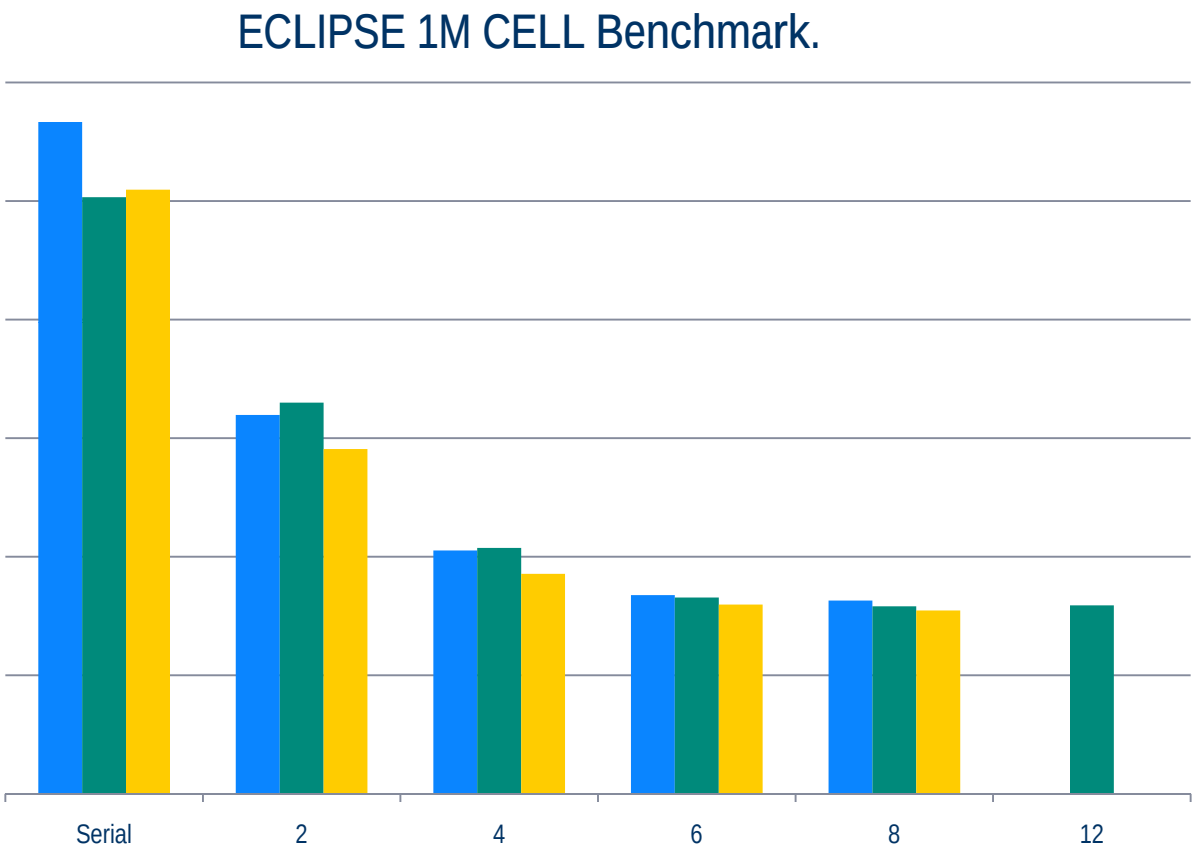

E300 Scaling

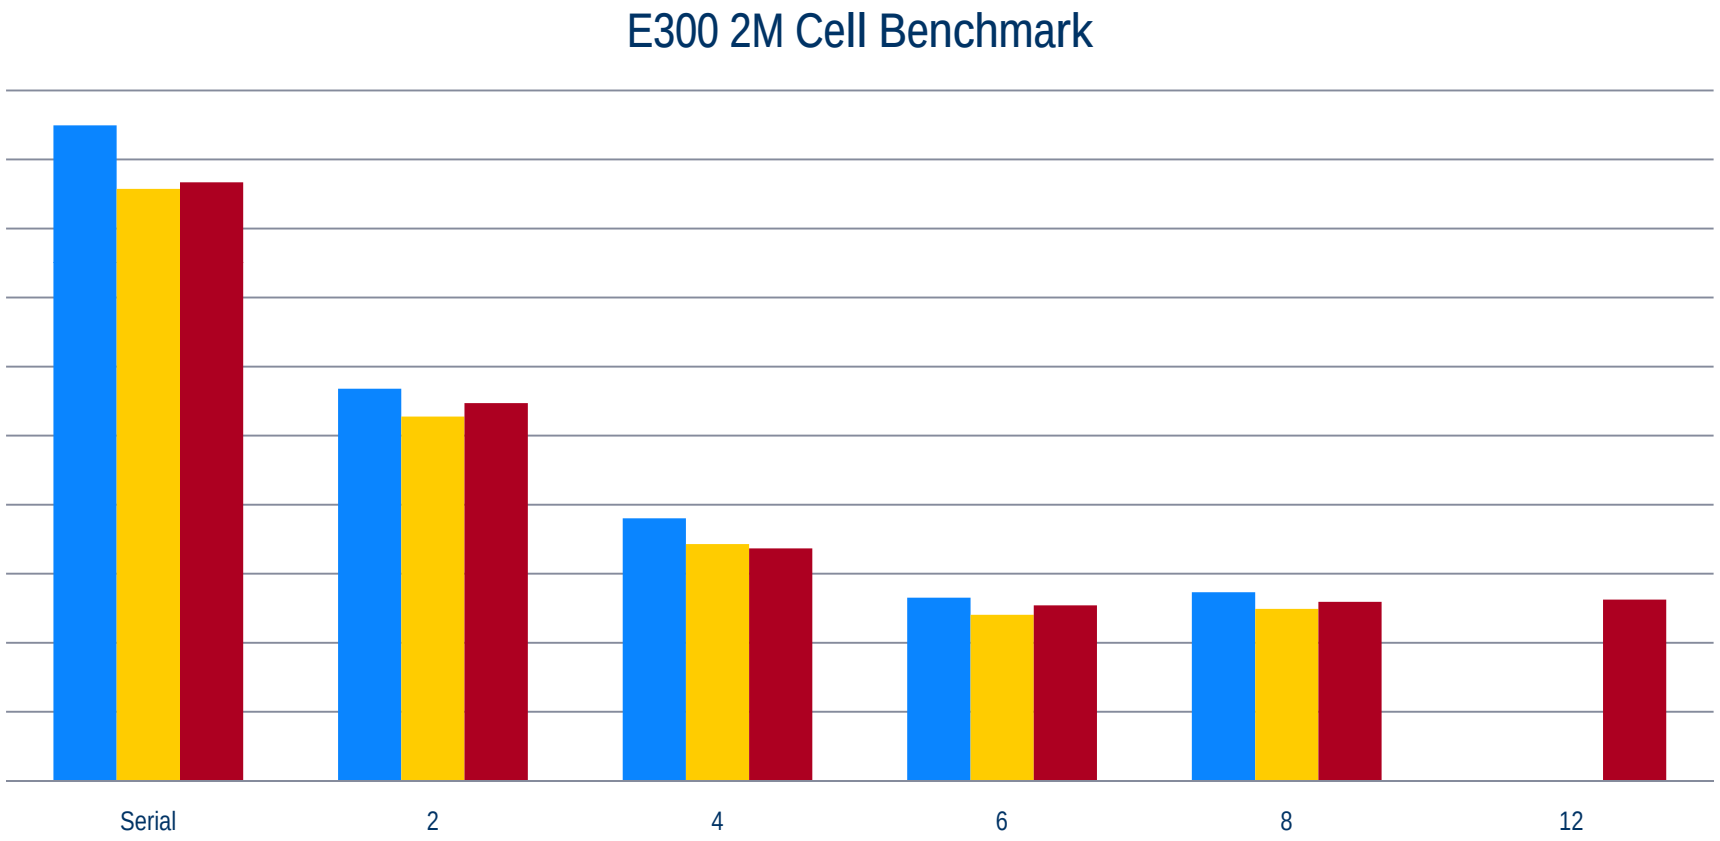

ECLIPSE scaling

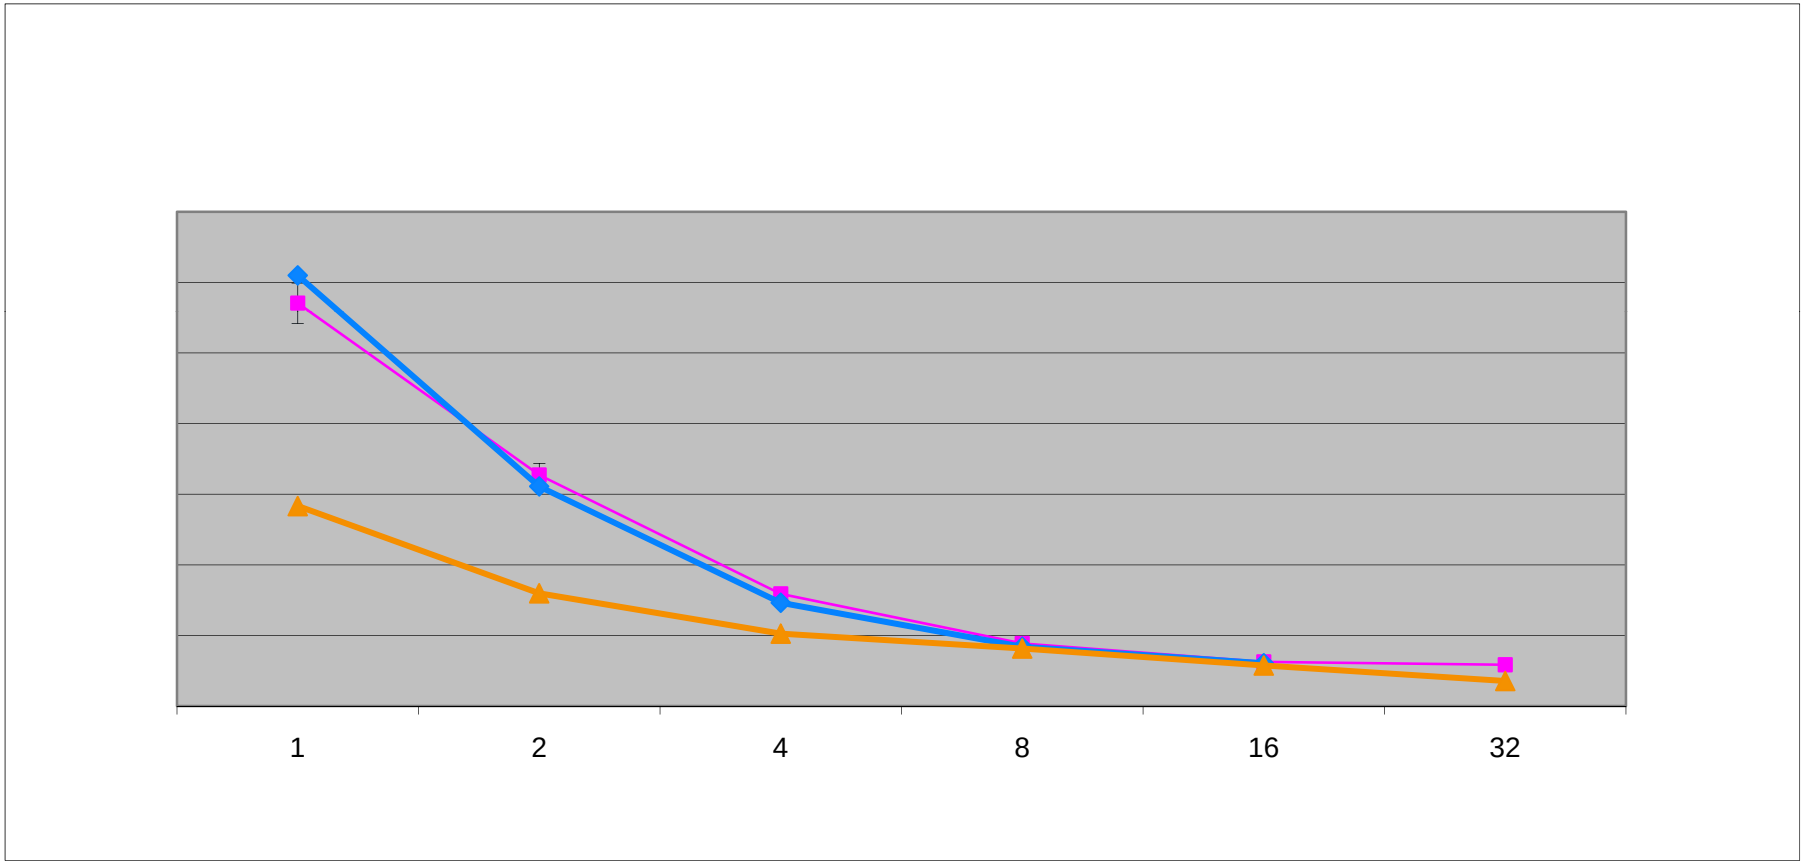

E300 Scaling

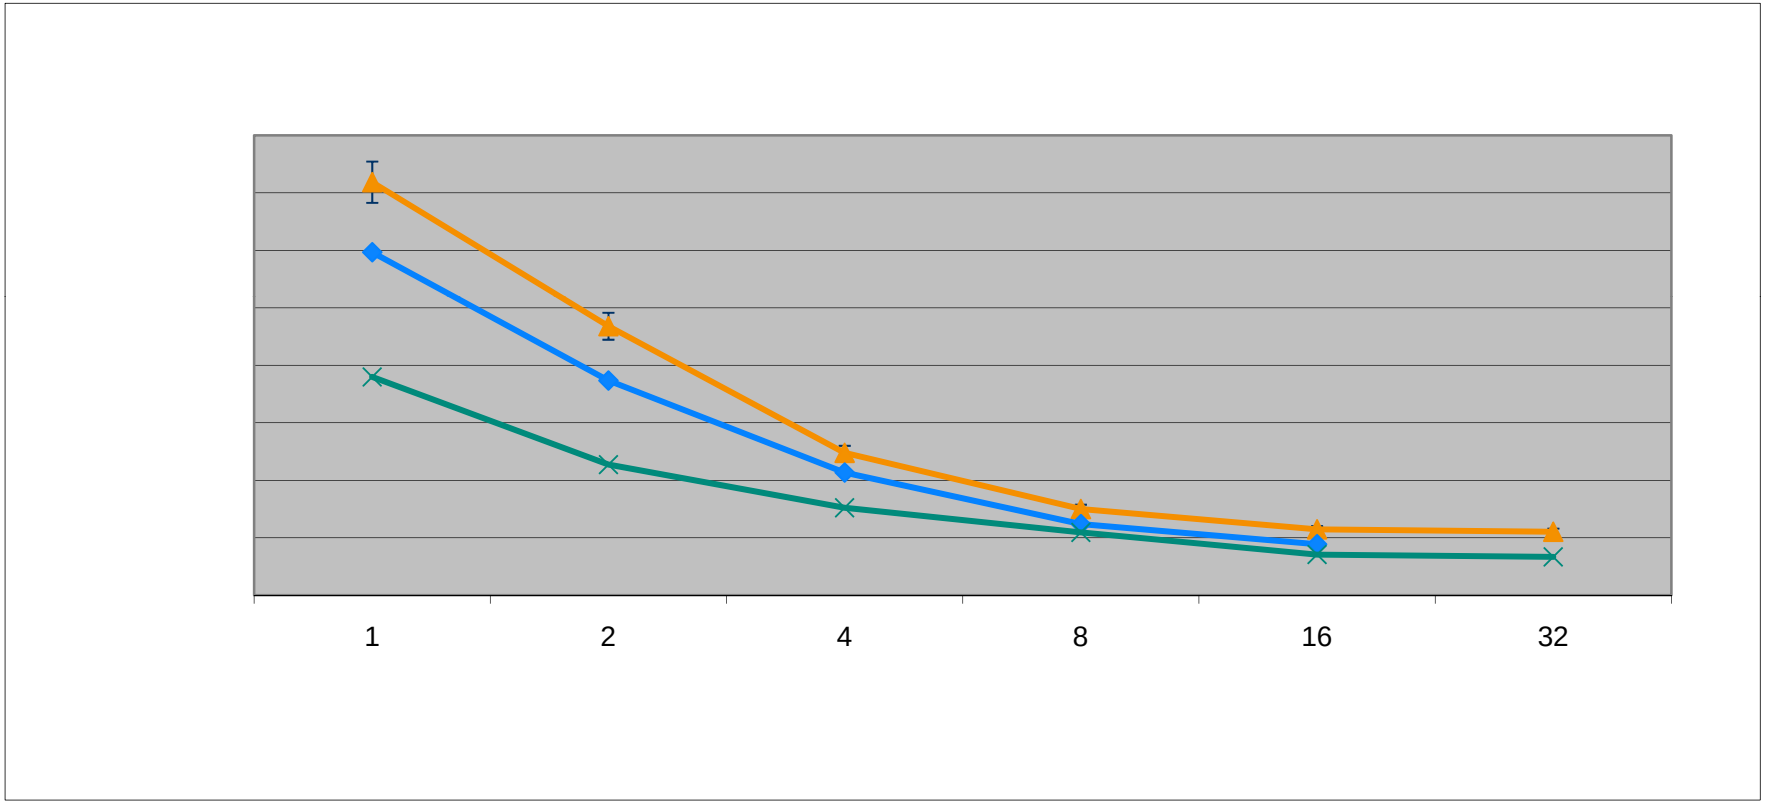

INTERSECT Scaling

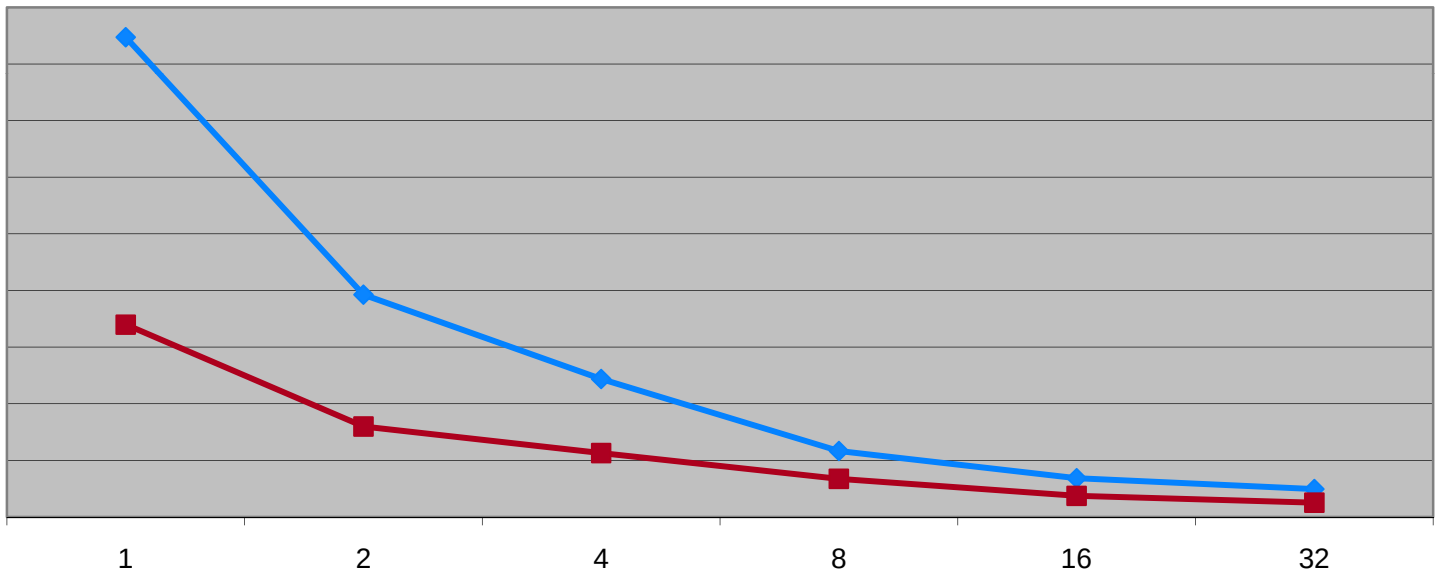

FrontSim scaling

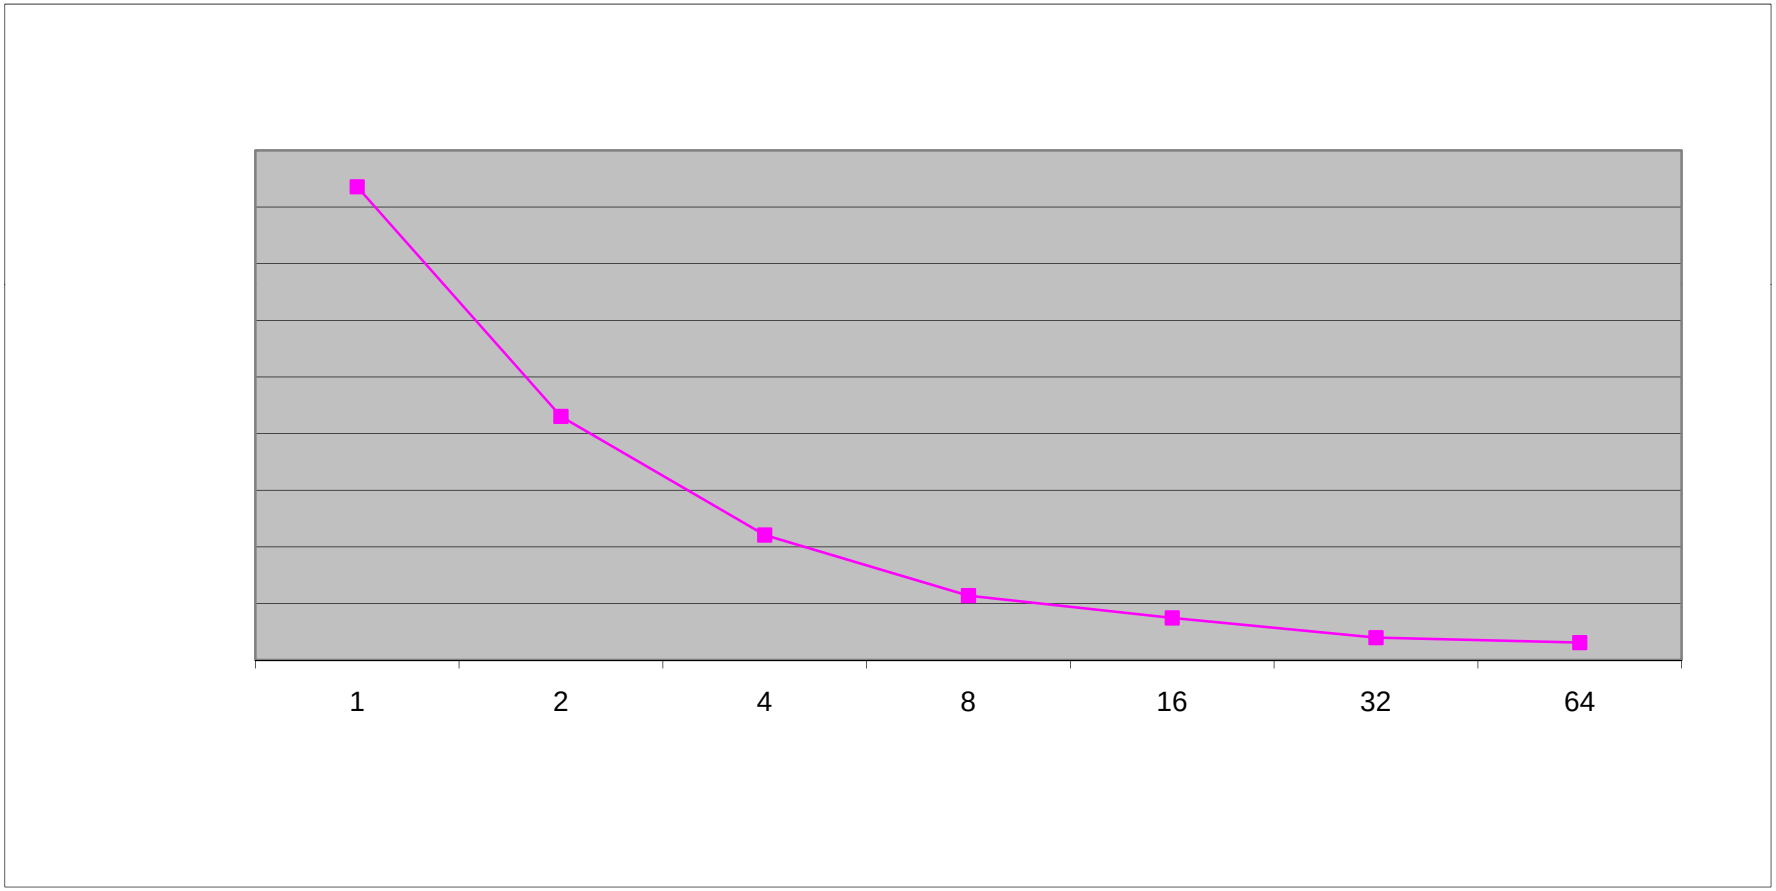

FrontSim scaling

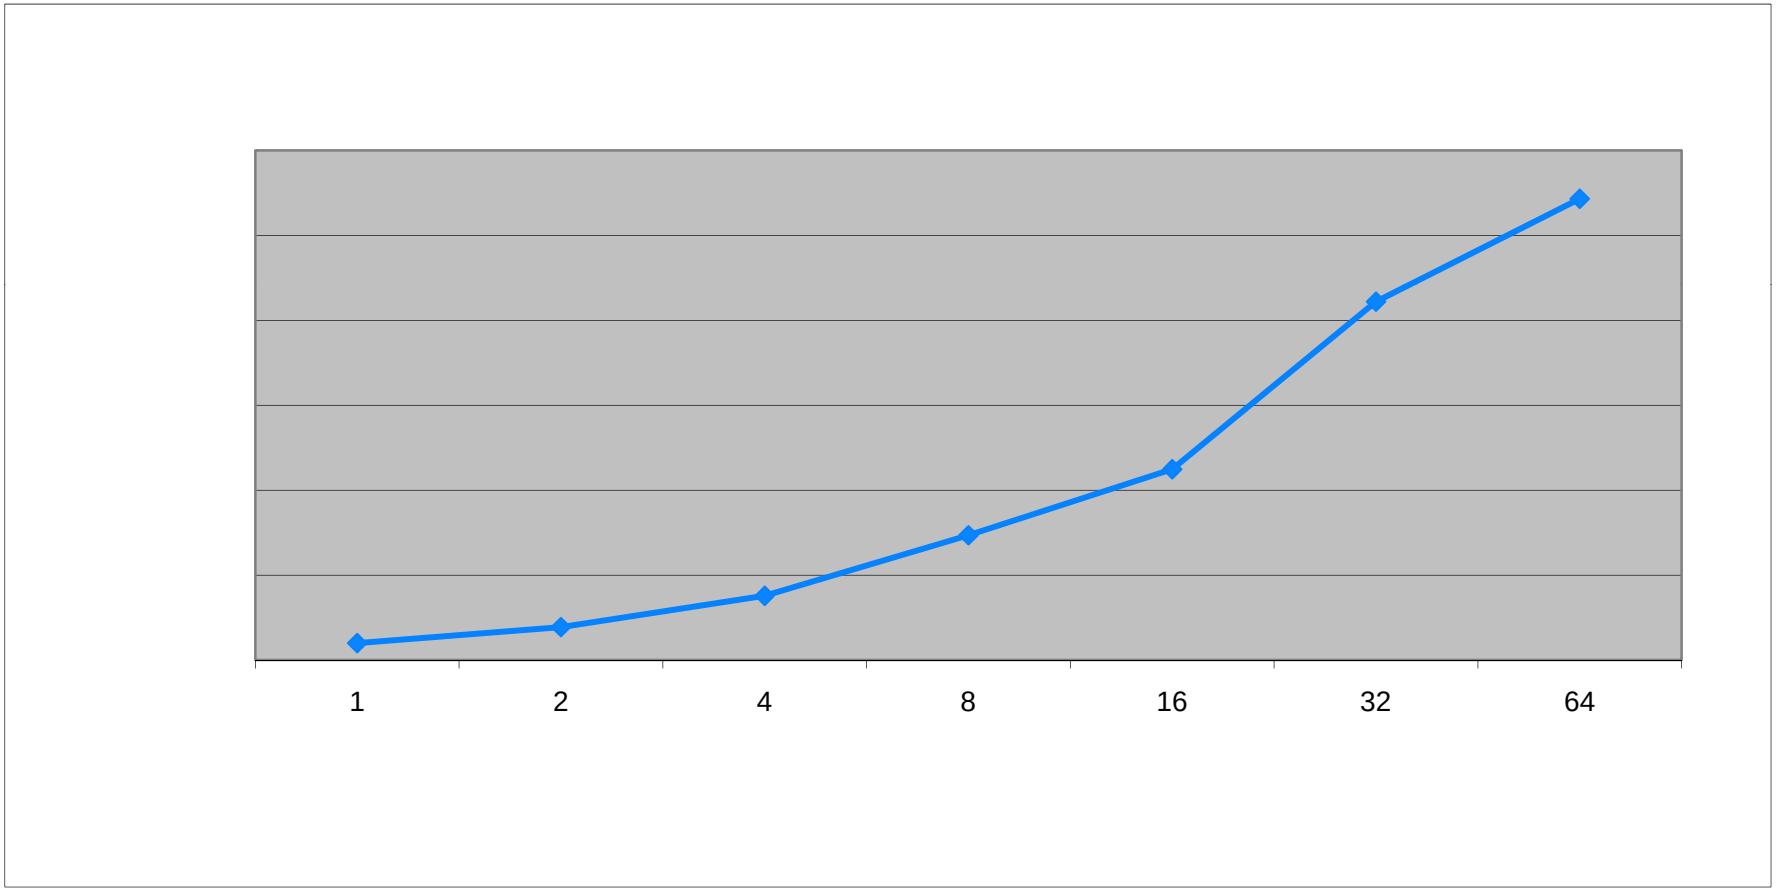

Multi-core chips interconnect becomes important

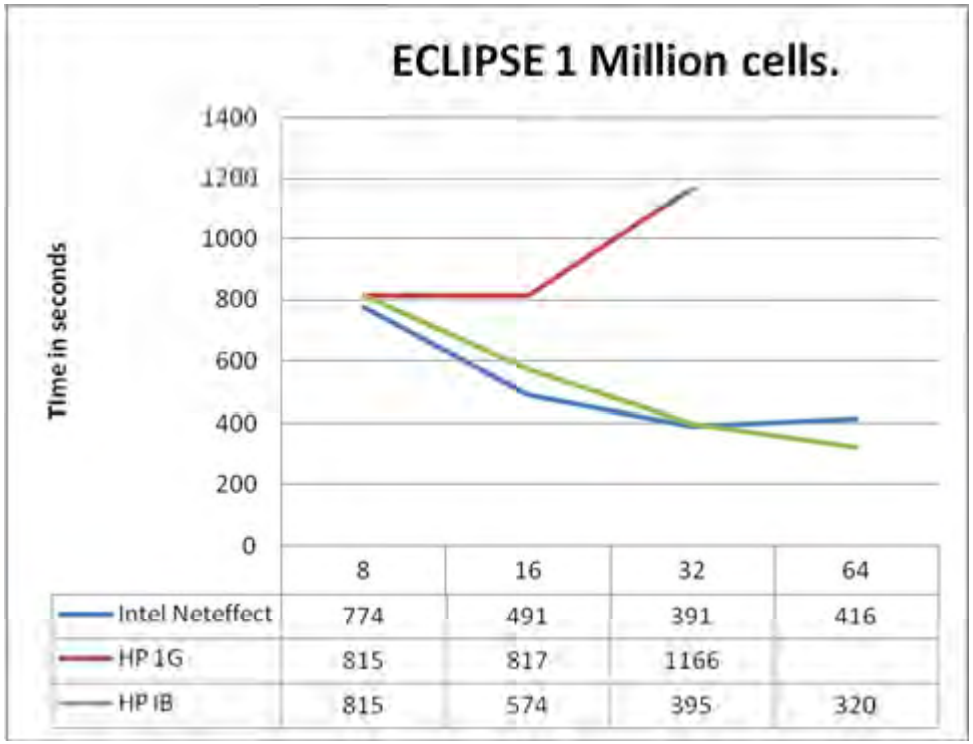

	8	16	32	64
Intel Neteffect	774	491	391	416
HP 1G	815	817	1166	
HP IB	815	574	395	320

Large number of core scaling

•Lot of misunderstandings about the scalability of the existing simulators.

- Model size and complexity are important when testing scalability
- Parallel scalability limited by the usual things
- Only real option is to try to run the model
  - Ideal scalability for E100 is 50000 cells per domain
  - Ideal scalability for E300 is 20000 cells per domain

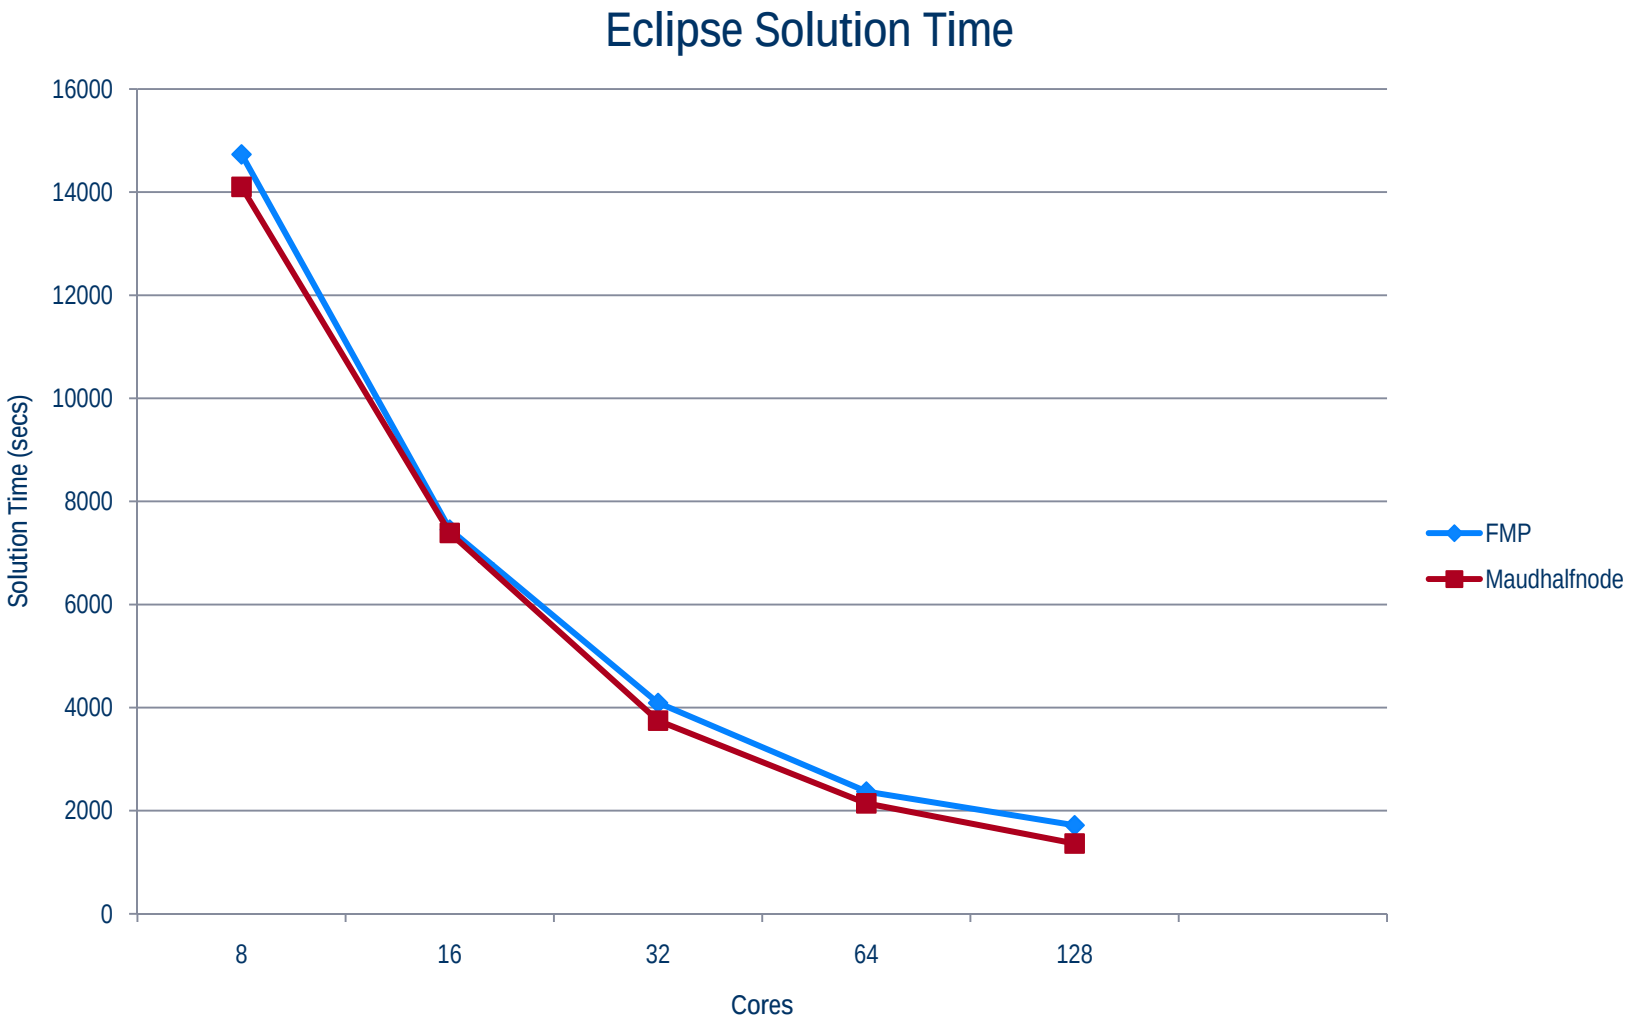
Eclipse Solution Time
16000
14000
12000
10000
8000
6000
4000
2000
0
Solution Time (secs)
8
16
32
64
128
Cores
FMP
Maudhalfnode


Schlumberger

E300 Solution Time (1.7 million cells compositional)

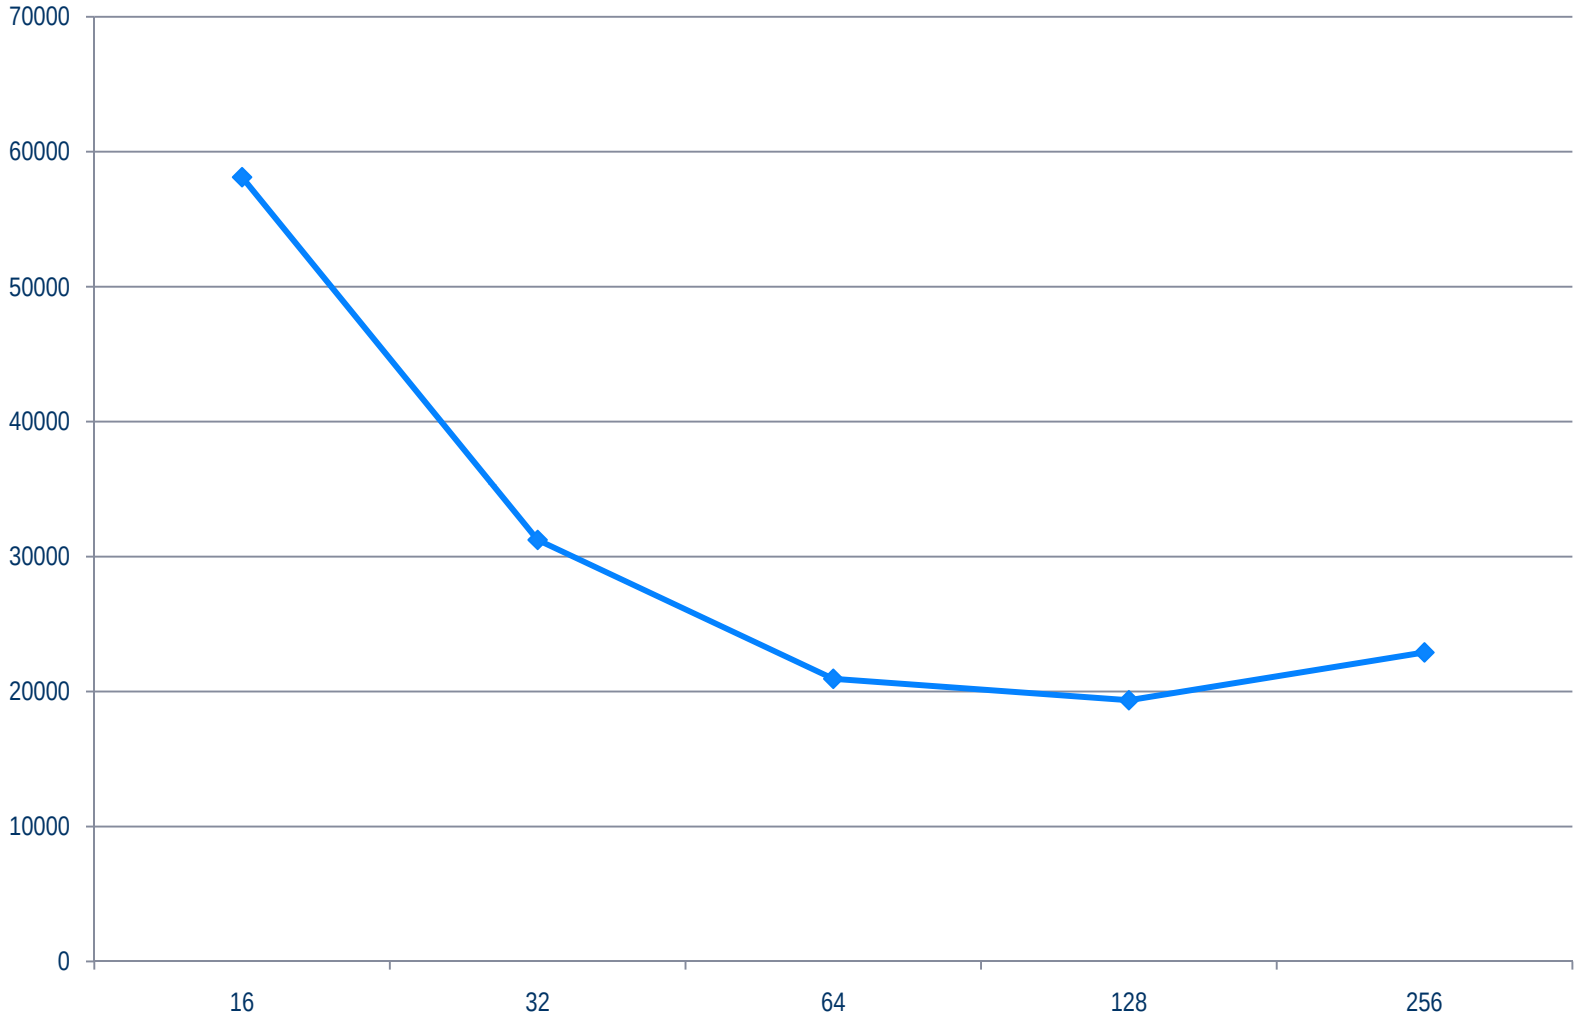

B

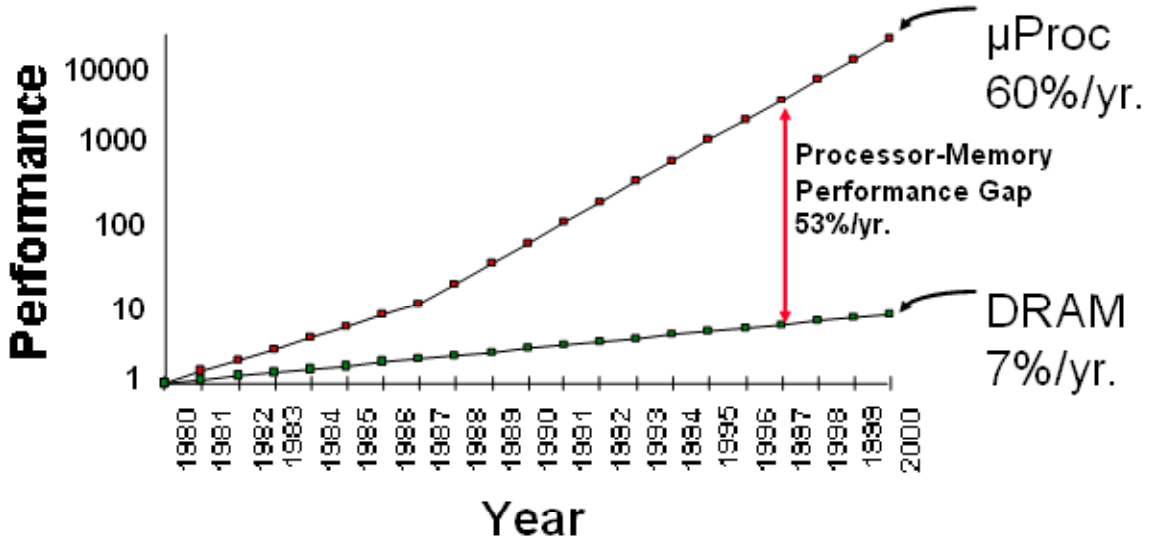

The CPU performance has increased 60% per year for almost two decades. On the other hand, the access time for main memory consisting of DRAM only decreased by 7-10% per year during the same period. The graph shown is excerpted from a talk slide of Trishul Chilimbi.

Summary

- Increasing cores/socket is not increasing speed or scalability
- Increasing cores/socket without an increase in memory bandwidth is not a lot of use for the distributed codes
- Multi-threaded codes such as FrontSim do have a significant benefit on the newer architectures
- Multi-core is the future and software developers will need to recode for this if they are to make use of this architecture
- Continued need for collaboration between the ISV's and the architecture vendors

INTERSECT parallel processing performance on a supergiant oil field

Paul A. Fjerstad

Project INTERSECT
Productivity and Innovation

- Project Intersect is a multiyear collaborative research project, involving Chevron and Schlumberger, that is aimed at developing next generation reservoir simulation software
- Researching into technology for large parallel unstructured simulation
- Already demonstrated new solver promising speed up ~3 in conventional simulation

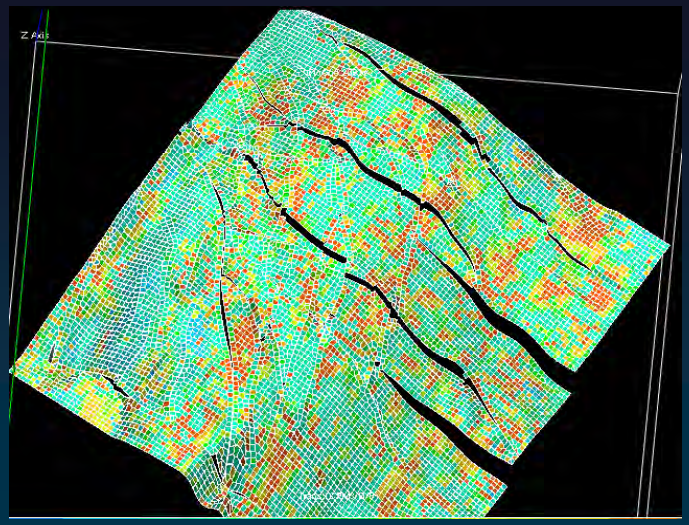

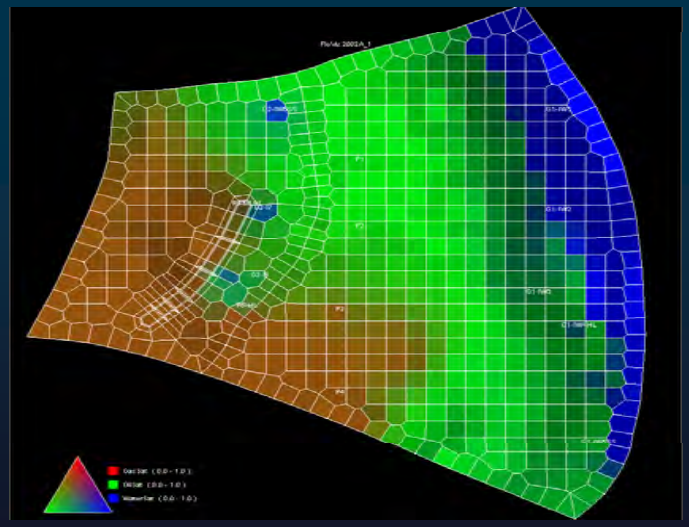

Speed, Large and Complex Models

Case study: A supergiant oilfield

Field characteristics:

- 6th largest oilfield in the world
- Partially fractured limestone lithology
- Sour gas (H_2S) miscible gas injection

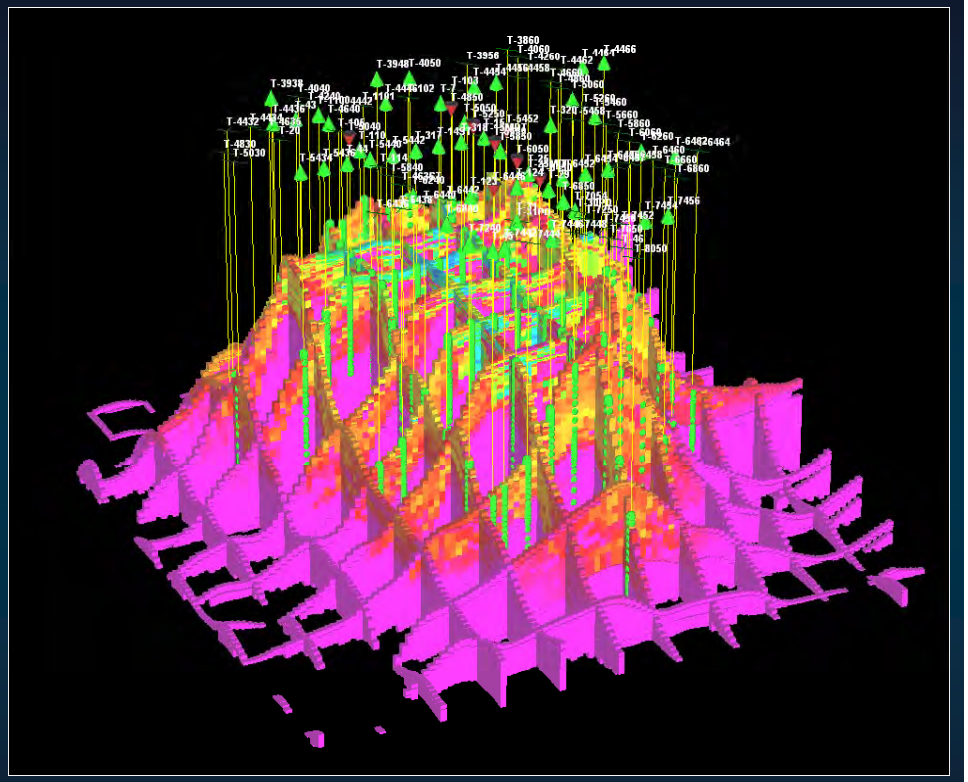

Fractured Reservoirs

Dual Porosity Dual Permeability

- Discontinuous series of sources and sinks surrounded by continuous fracture medium
- Fluid flow exchange between fracture and matrix
- New phase behaviour approach facilitates simulation of large compositional models

Wellbore image showing fractures

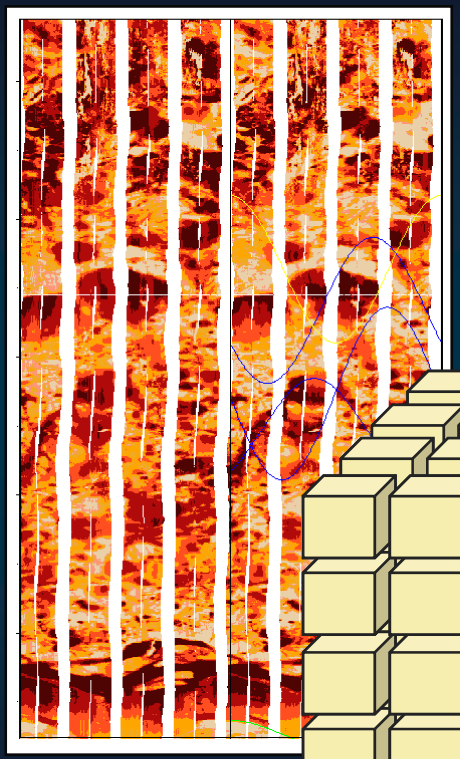

Numerical representation

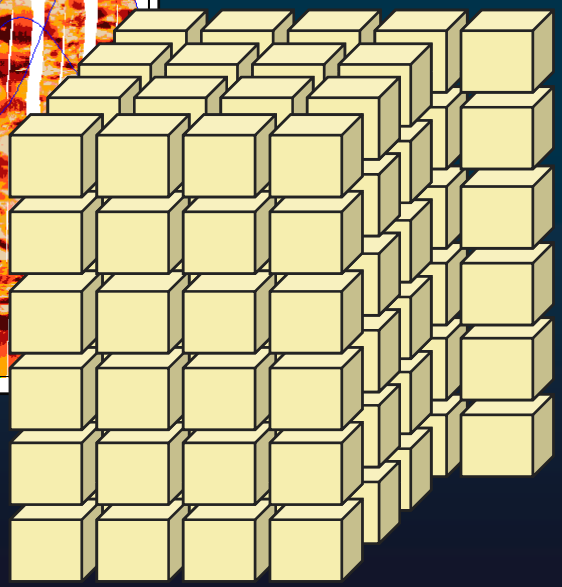

Multilevel Graph Partitioning

- Computational load modeling with node weights
- Computational communication cost modeling using connection weights
  - minimize communication across processors and avoid partitioning at nodes with lots of connections!
  - transmissibility
- Static partitioning computed *a priori*

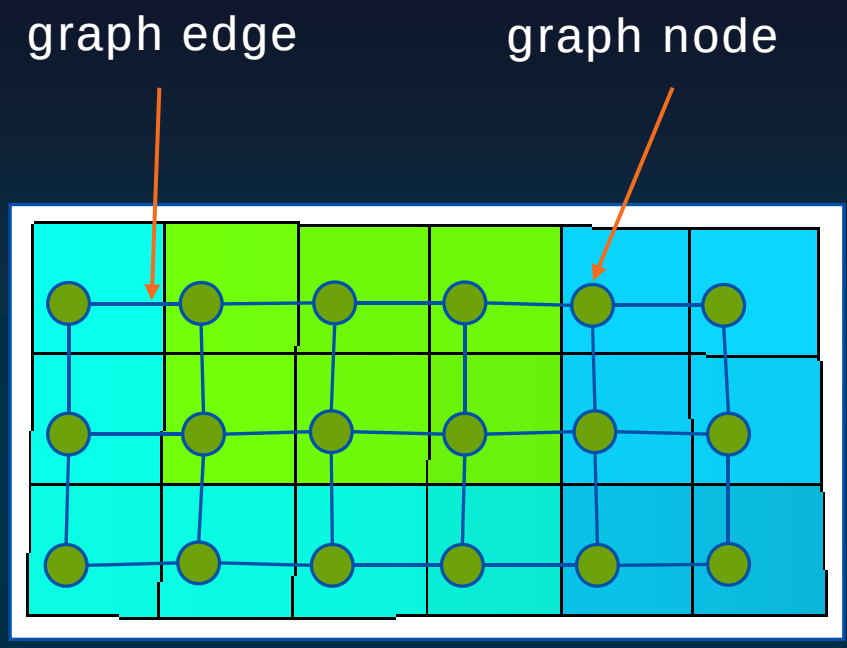

3 domains with 6 nodes on each processor

Cluster hardware comparisons

- Legacy cluster is 2008 configuration
- 2010 cluster is on average 2.25 x faster than legacy cluster

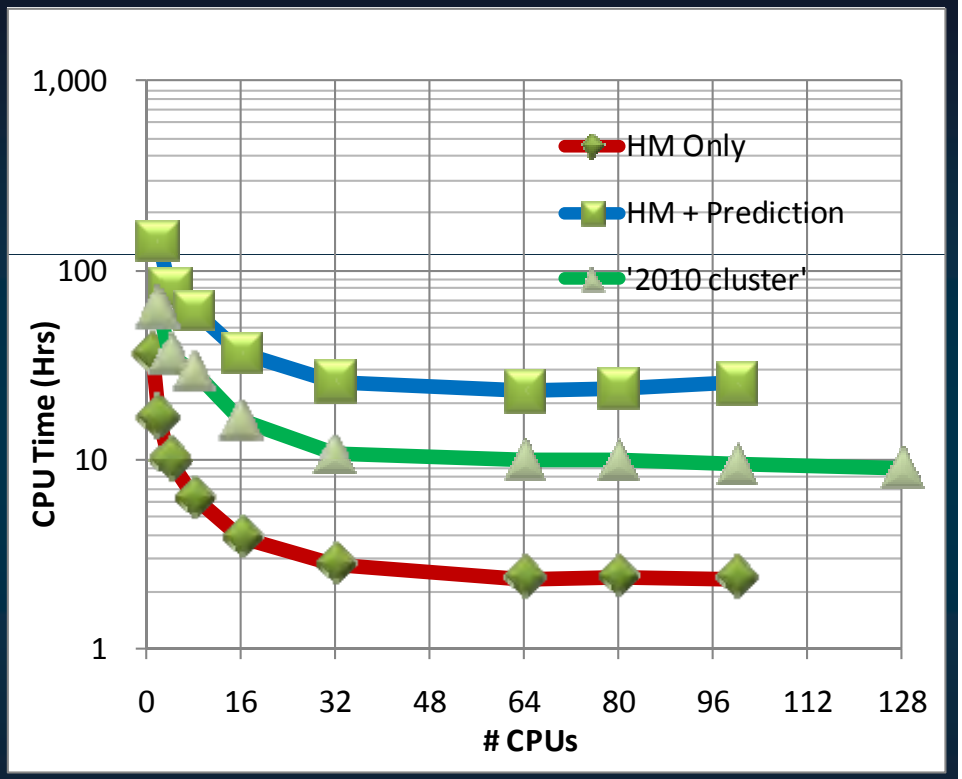

Parallel Scalability

Computational overhead

- Deviate from optimum scalability
- Increase in cpu time
- Communication and synchronization
- Non-parallel (serial) part of the program
  - (one processor works, others spin idle)

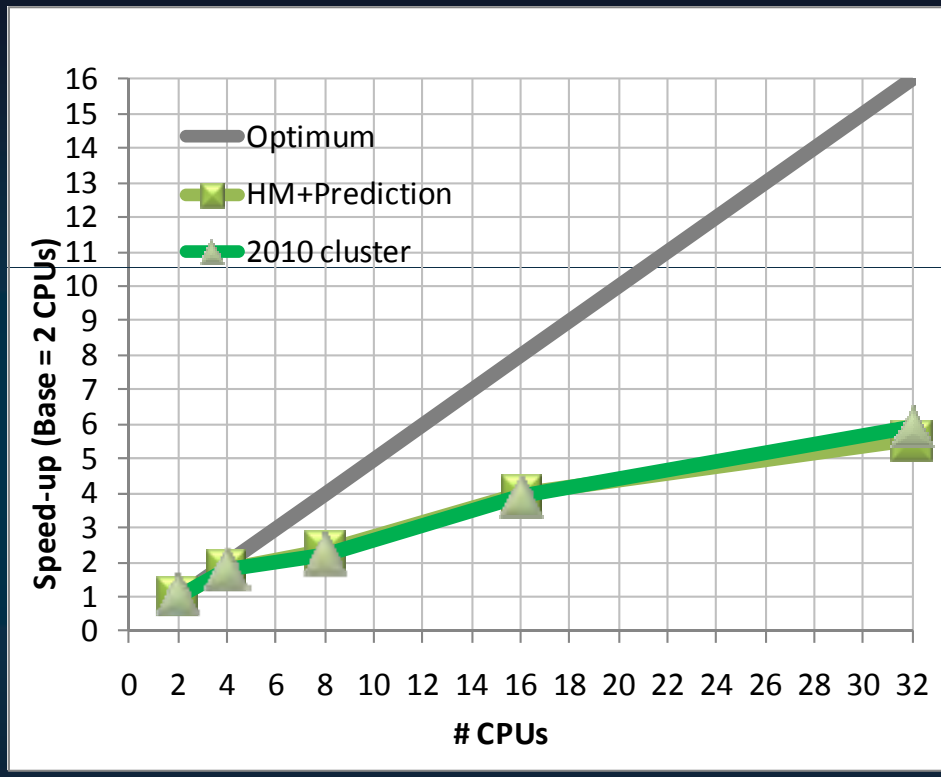

Case study: A supergiant oilfield

Business benefits summary:

- More realistic geological descriptions to give more accurate production forecasts
- Capability to include more realistic fluid descriptions
- Reduced runtimes, allowing better assessments of uncertainty providing a basis for better decision making