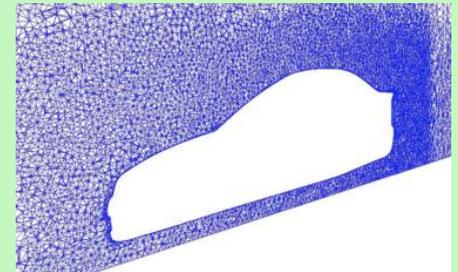
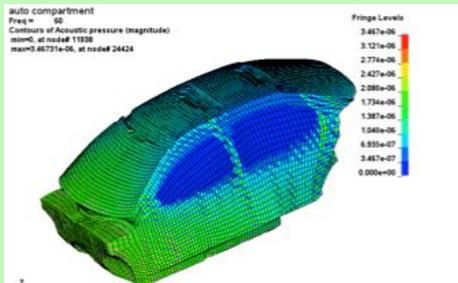


Productivity and Sustainability in Disruptive Times



Bob Lucas
Oct. 16, 2015



Outline

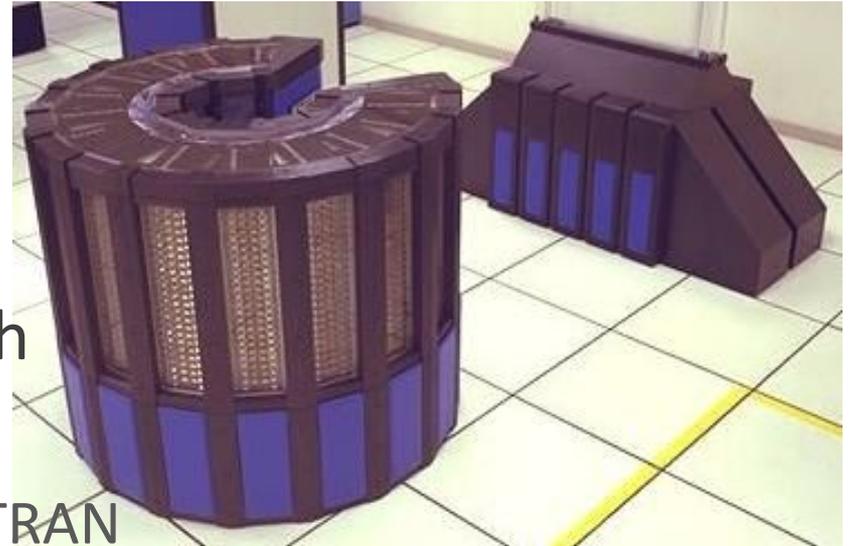
- CSE requirements are unbounded
- How we got to where we are today
- How LS-DYNA is evolving
- Thoughts about the future

Computational Requirements are Unbounded

- The only bound on the CSE demand for more sophisticated algorithms, running on bigger and faster computing systems, is human imagination
- Applications span national security, scientific discovery, industrial competitiveness, commerce, public health, and yes, entertainment
DreamWorks thinks of itself as a manufacturing business
- Mechanical Computer Aided Engineering (MCAE)
LSTC users could easily consume orders-of-magnitude more computational capability

Productivity – Old School

- When I started my career, machines were expensive, people were cheap
- Linpack correlated well with MCAE codes
“Floating point engines for FORTRAN applications”
- Increasing productivity meant increasing floating point operations per second (Flop/s) per dollar



NASA Cray-2 from Wikipedia

Disruptive Change – Circa 1990

- There was a lot of uncertainty and excitement about the future in the late 1980s and early 1990s

CMOS replacing bipolar

Lots of imaginative architectures

Lots of alternative execution models



CM-5



MasPar



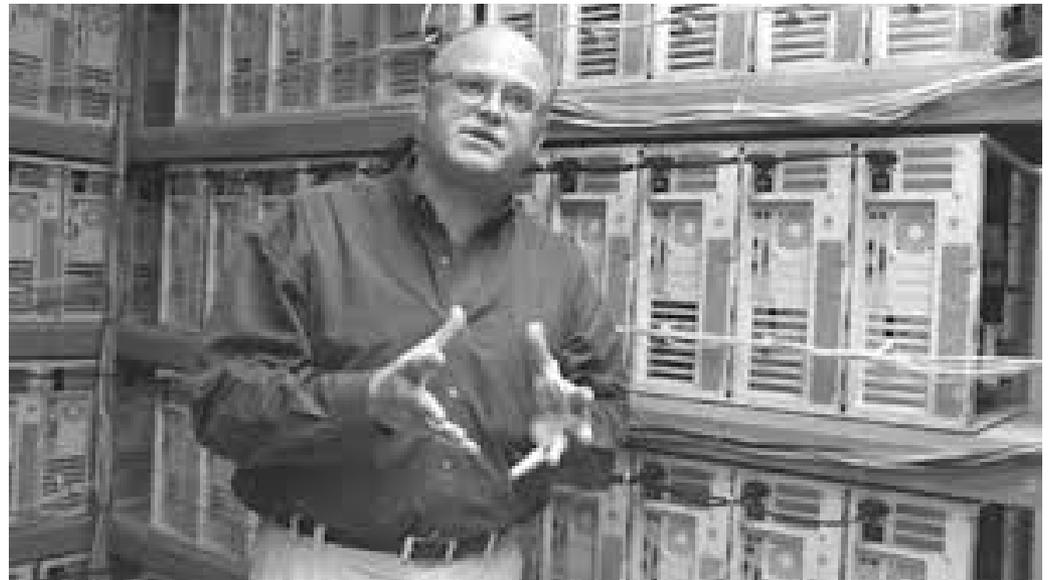
Convex SPP



Cray T3E

Today's Model Emerged

- Winning model was NASA GSFC's Beowulf
Communicating sequential processes execution model
- Cheapest possible computing system
Commodity H/W (PC components)
Free S/W (Linux)
- “Dark Matter”



Thomas Sterling and Caltech Beowulf

Unanticipated Side-Effect

- Fragmenting the address space was very disruptive
- Bad for productivity and sustainability
Programmers reason about data placement and choreography
Constrained by functions in libraries (IEEE 1516 HLA/RTI)
- Good news is that often most of a code was unchanged
Sandia's James Peery says 90%
- Many codes could evolve
But many others couldn't

False Economy

- H/W is cheap
Buy it by the acre



- Use it inefficiently (Hadoop)
Who cares, computing is “too cheap to meter”
- But programming is human labor
MPI applications are expensive software
Increases the cost of science and engineering

Facebook Prineville Data Center

Threatening to Happen Again

- Increasing diversity of general purpose architecture
 - DOE community speaks of two “swim lanes”
 - Intel alone has multiple products for CSE
 - AMD APUs and NVIDIA GPU accelerators
 - Memories too: HMC, HBM and Xpoint
- Specialized processor products
 - IBM neuromorphic processor
 - Micron automata processor
- Could get a lot more diverse
 - IARPA superconducting program
 - Quantum annealing
 - Maybe even quantum computing



Will This Lead to a Software Crisis?

- Frog in the Pot metaphor

Its been slowly getting warmer for decades

- Possible software affordability crisis

What if we can't afford to keep pace?

I.e., productivity does not keep pace with growth in complexity

- Possible software capability crisis

What if problems can't be solved, regardless of effort?

“General purpose” machines run a small subset of our problems

How can you use specialized machines when your algorithms are themselves subjects research

LS-DYNA Genesis

- Writing a sophisticated multi-physics code takes a very long time

LLNL started DYNA3D in 1976 for the Cray 1
John Hallquist formed LSTC in 1989

- LS-DYNA is not an old code

NASTRAN started in 1964 (1st release in '68)
ANSYS founded in 1970

- Risky endeavor

Many attempts at large new software projects fail

Post and Kendall, The International Journal of High Performance Computing Applications, Volume 18, No. 4, Winter 2004

Example of What LS-DYNA is Used For



LS-DYNA Has Evolved over the Years

- Distributed memory computers

New hire in 1993 to begin distributed memory (pre MPI)

We need to run on the same platforms as everybody else

Large government purchases shape the HPC environment

- New engineering features

Subsumed NIKE3D to add implicit time integration

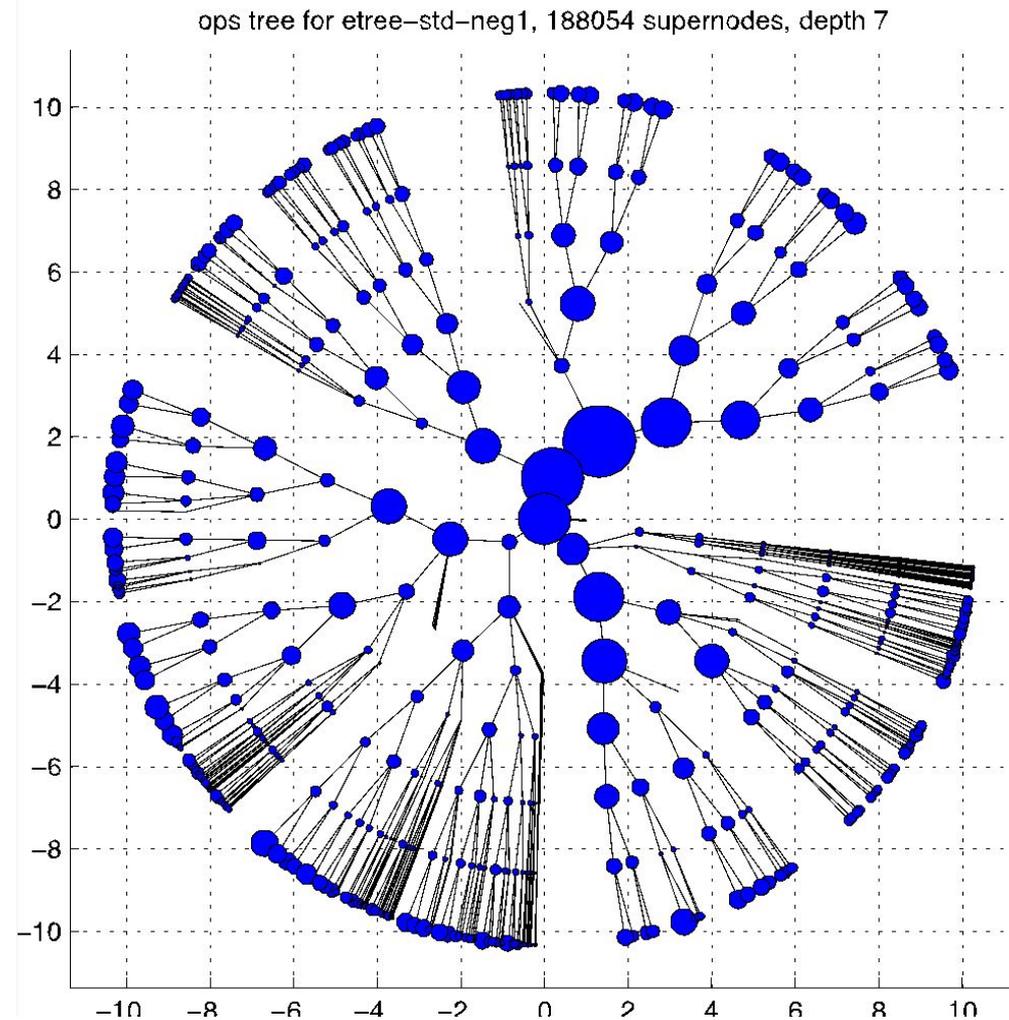
Key to metal forming

- New mathematical algorithms

Hired me to get a scalable linear solver

Evolution is Not Easy

- Hard scaling problems have to be overcome
 - Reordering is NP-complete
 - Load-balancing is NP-hard
- What if the best heuristics are sequential?
 - Some of the biggest challenges to exascale are going to be mathematical



Chevy Silverado elimination tree
Ordered by Metis, weighted by Flops

LS-DYNA Today

- O(10M) source lines of code (SLOCs)

Growth is accelerating (more people)

- Primarily Fortran, ~10% C

MPI for distributed memory

OpenMP (and the like) for multi-core nodes

A bit of CUDA for GPU accelerators

- Runs on tens of thousands of cores

Customers want more

“Also the current code is limited to 4096 processes so I cannot run the job up to the 96k cores I wanted to.”

LS-DYNA Will Evolve Forward

- We won't rewrite LS-DYNA from scratch
It would not be a productive use of time or money
- More importantly, we wouldn't want to
It works
Its developers understand it (at least their subsets)
Its MCAE users trust it
Intel says it "moves the most Pentiums" in MCAE
- More value in S/W today than in H/W
When S/W doesn't port, machines risk becoming irrelevant
Already happened, SPICE and EEs on Cray vector machines

Computational Research and Engineering Acquisition Tools and Environments (CREATE)

CREATE is a multi-phase program that started in 2008, to develop and deploy four (now five) computational engineering tool sets for acquisition engineers



- ❑ **Aircraft (AV) Design Tools:** Fixed-wing aircraft, rotorcraft, conceptual design, trade-space exploration and operational testing and transition
- ❑ **Ship Design Tools:** Shock/damage, hydrodynamics, early-stage design & trade-space exploration, and operational testing and transition
- ❑ **Radio Frequency (RF) Antenna Design and Integration Tools:** Conceptual design and detailed analysis tools relevant to virtually all DOD platforms
- ❑ **Ground Vehicles (GV) Tools:** End-to-end mobility solver, provide rapid, physics-based data for design and trade-space analysis
- ❑ **Meshing and Geometry (MG) Support:** The geometry and meshing project improves the ease, speed, flexibility, and quality of geometry and mesh generation, and enables the generation of CAD-neutral digital representations and product models of weapons systems & platforms and operational terrains and environments



CSE Needs Revolution Too, to Maximize Innovation

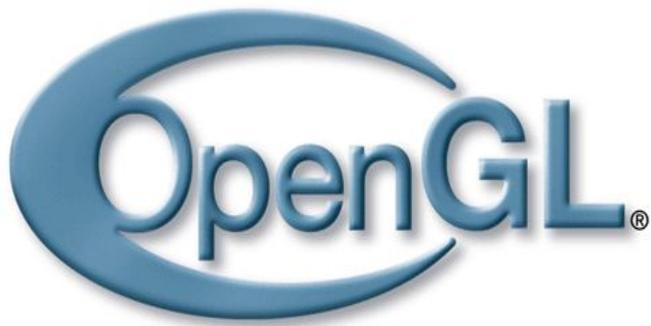
- DYNA3D was a response to disruptive technology
Cray 1 vectors
- Implicit MCAE codes evolved onto Crays
Their computational bottlenecks (linear solvers) vectorized
- Innovation requires investment
Academe (UIUC ILLIAC-IV & Caltech Cosmic Cube)
National Labs (LLNL DYNA3D, GSFC Beowulf)
Industrial labs (SUN's Java, IBM's 801)
DOD CREATE
- Successful innovations will be assimilated

Take Control of Our Own CSE Destiny

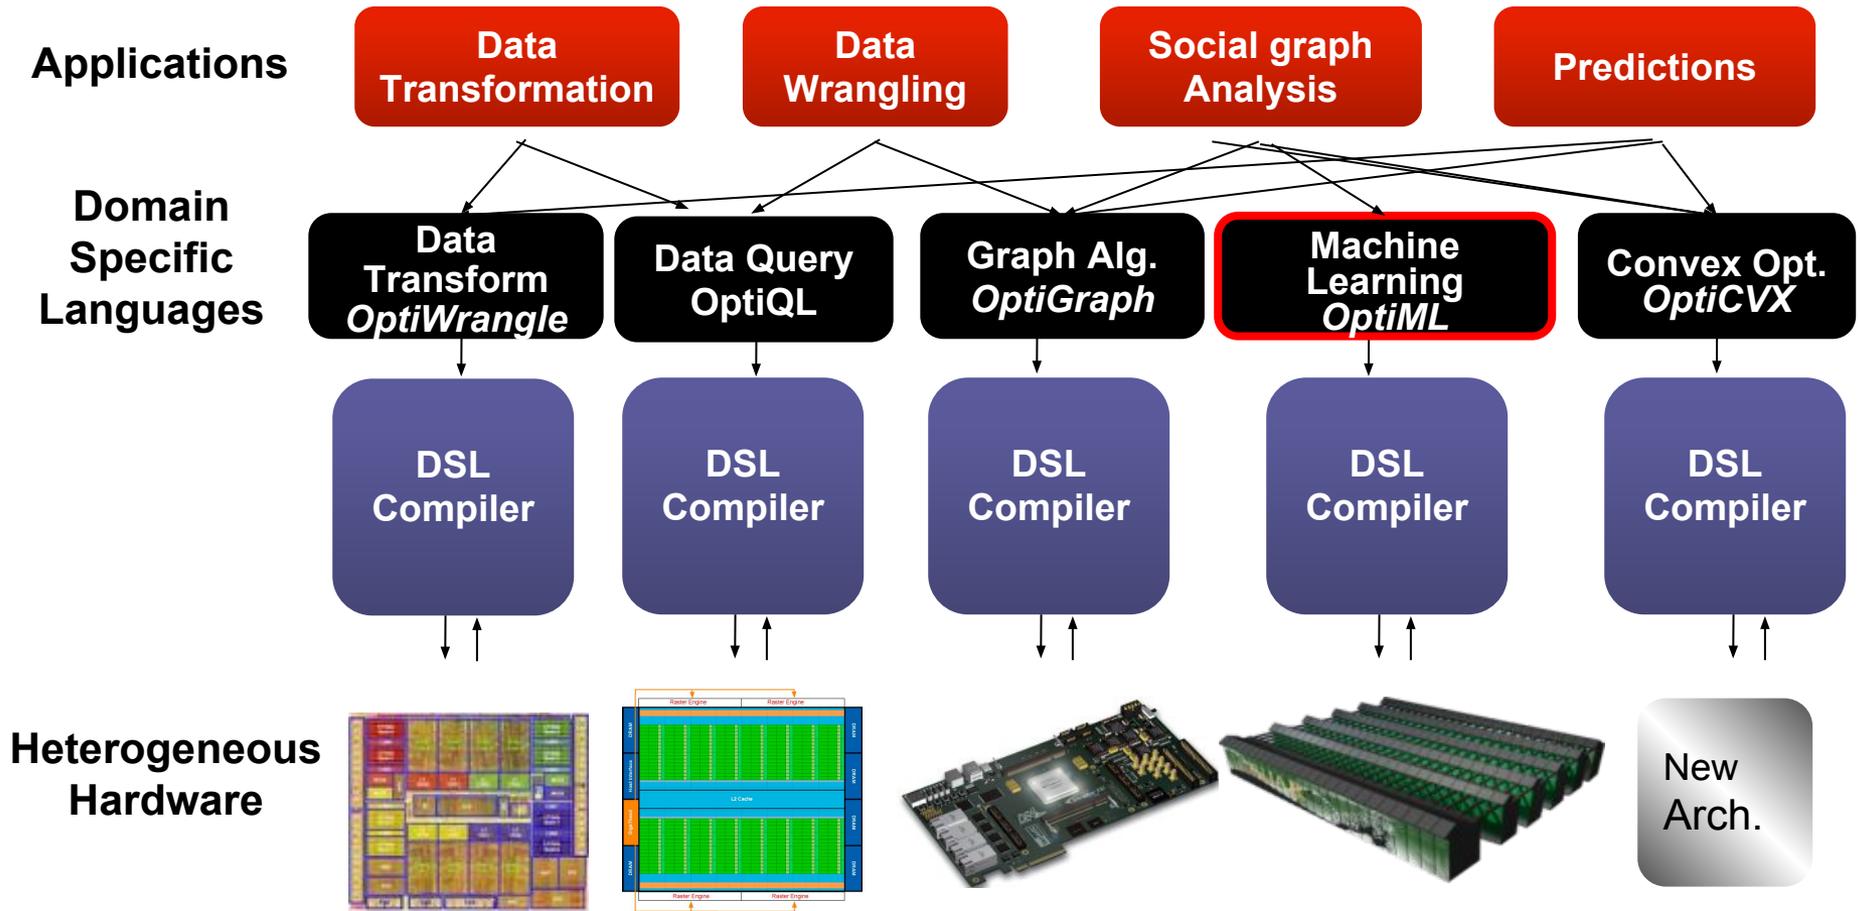
- To maintain, much less improve productivity and sustainability, we need to take control of our destiny
 - Maximize CSE throughput, not Flop/s per dollar
 - Blue Waters not on the TOP 500
- New COTS ecosystem is the mobile SOC
 - SOC ecosystem allows for specialization
 - We should embrace this opportunity today
 - We embraced PCs 20 years ago
- Why not families of specialized machines?
 - Anton (100X capability)
 - Green Flash (100X energy efficiency)
 - HP's "Bespoke SOC"

Domain Specific Languages (Not Just H/W)

- A language (or library) that exploits domain knowledge for productivity and efficiency
High-level, usually declarative, and deterministic



Stanford's Delite DSL Framework



Courtesy Kunle Olukotan and Pat Hanrahan

Sustainability Requires Investment Too

- Need to invest in S/W maintenance and enhancement
Just like physical infrastructure and scientific instruments
Geant4 is as much a part of experimental physics as the instruments it models
- LS-DYNA uses BCSLIB-EXT, LAPACK, Metis, MPI, MUMPS, SVN, etc.
Who ensures support if their authors are on soft money?
- Nothing new, Ken Kennedy spoke about this problem twenty years ago.

Summary

- CSE demand is unbounded
- False economy of today's systems
Labor is expensive
- Software like LS-DYNA must be able to evolve
But there must also be room for innovation
- We need to take control of our destiny
Focus on human productivity this time around



Thank you!