

Hey Buddy, Can You Spare \$1.50? Or Neat Hardware, Now Whatcha Gonna Do With It?

David E. Bernholdt

Computer Science and Mathematics
Oak Ridge National Laboratory

bernholdtde@ornl.gov

(virtual nobody)

Oak Ridge National Laboratory is managed by UT-Battelle, LLC
for the US Dept. of Energy under contract DE-AC-05-00OR22725.

The Hardware Is Getting Interesting Again

- Clusters are not *the* solution they once were
 - Al G's comment about SOS6
- SSI and DSM are back with a vengeance
- Ability to configure hardware "to order"
- "Dynamic range" of architectures is huge
 - Vector, commodity CPU, reconfigurable co-processor
 - Different balances (CPU, memory, interconnect)
 - Steve R, John L, Manish G, others
 - Order of magnitude difference in # processors required to reach 100 TFLOPS (4,000 vs 40,000)
 - Jeff N

Most What ^ Users Want

- They want to do their science
 - Having to worry about “computer science” is a distraction

- Write once, run anywhere
 - Platform *availability* much more important than architecture
 - Platforms change rapidly
 - Performance portability

The Price of Success

- Successful applications scientists are asked to do larger, high-fidelity simulations
- Software gets more complex to handle larger, more realistic problems
- Successful hardware vendors are asked to produce faster and faster computers
- Hardware gets more complex to provide performance
- Software gets even more complex because hardware complexity must be managed at the user's level
- Comments by Jeff N, Juan M

Example from Quantum Chemistry: Tensor Contraction Engine

$$\begin{aligned} \text{hbar}[a,b,i,j] = & \text{sum}[f[b,c]*t[i,j,a,c],\{c\}] - \text{sum}[f[k,c]*t[k,b]*t[i,j,a,c],\{k,c\}] + \text{sum}[f[a,c]*t[i,j,c,b],\{c\}] - \text{sum}[f[k,c]*t[k,a]*t[i,j,c,b],\{k,c\}] - \\ & \text{sum}[f[k,j]*t[i,k,a,b],\{k\}] - \text{sum}[f[k,c]*t[j,c]*t[i,k,a,b],\{k,c\}] - \text{sum}[f[k,i]*t[j,k,b,a],\{k\}] - \text{sum}[f[k,c]*t[i,c]*t[j,k,b,a],\{k,c\}] \\ & + \text{sum}[t[i,c]*t[j,d]*v[a,b,c,d],\{c,d\}] + \text{sum}[t[i,j,c,d]*v[a,b,c,d],\{c,d\}] + \text{sum}[t[j,c]*v[a,b,i,c],\{c\}] - \text{sum}[t[k,b]*v[a,k,i,j],\{k\}] + \text{sum}[t[i,c]*v[b,a,j,c],\{c\}] - \\ & \text{sum}[t[k,a]*v[b,k,j,i],\{k\}] - \text{sum}[t[k,d]*t[i,j,c,b]*v[k,a,c,d],\{k,c,d\}] - \text{sum}[t[i,c]*t[j,k,b,d]*v[k,a,c,d],\{k,c,d\}] - \text{sum}[t[j,c]*t[k,b]*v[k,a,c,i],\{k,c\}] \\ & + 2*\text{sum}[t[j,k,b,c]*v[k,a,c,i],\{k,c\}] - \text{sum}[t[j,k,c,b]*v[k,a,c,i],\{k,c\}] - \text{sum}[t[i,c]*t[j,d]*t[k,b]*v[k,a,d,c],\{k,c,d\}] \\ & + 2*\text{sum}[t[k,d]*t[i,j,c,b]*v[k,a,d,c],\{k,c,d\}] - \text{sum}[t[k,b]*t[i,j,c,d]*v[k,a,d,c],\{k,c,d\}] - \text{sum}[t[j,d]*t[i,k,c,b]*v[k,a,d,c],\{k,c,d\}] \\ & + 2*\text{sum}[t[i,c]*t[j,k,b,d]*v[k,a,d,c],\{k,c,d\}] - \text{sum}[t[i,c]*t[j,k,d,b]*v[k,a,d,c],\{k,c,d\}] - \text{sum}[t[j,k,b,c]*v[k,a,i,c],\{k,c\}] - \text{sum}[t[i,c]*t[k,b]*v[k,a,j,c],\{k,c\}] - \\ & \text{sum}[t[i,k,c,b]*v[k,a,j,c],\{k,c\}] - \text{sum}[t[i,c]*t[j,d]*t[k,a]*v[k,b,c,d],\{k,c,d\}] - \text{sum}[t[k,d]*t[i,j,a,c]*v[k,b,c,d],\{k,c,d\}] - \\ & \text{sum}[t[k,a]*t[i,j,c,d]*v[k,b,c,d],\{k,c,d\}] + 2*\text{sum}[t[j,d]*t[i,k,a,c]*v[k,b,c,d],\{k,c,d\}] - \text{sum}[t[j,d]*t[i,k,c,a]*v[k,b,c,d],\{k,c,d\}] - \\ & \text{sum}[t[i,c]*t[j,k,d,a]*v[k,b,c,d],\{k,c,d\}] - \text{sum}[t[i,c]*t[k,a]*v[k,b,c,j],\{k,c\}] + 2*\text{sum}[t[i,k,a,c]*v[k,b,c,j],\{k,c\}] - \text{sum}[t[i,k,c,a]*v[k,b,c,j],\{k,c\}] \\ & + 2*\text{sum}[t[k,d]*t[i,j,a,c]*v[k,b,d,c],\{k,c,d\}] - \text{sum}[t[j,d]*t[i,k,a,c]*v[k,b,d,c],\{k,c,d\}] - \text{sum}[t[j,c]*t[k,a]*v[k,b,i,c],\{k,c\}] - \text{sum}[t[j,k,c,a]*v[k,b,i,c],\{k,c\}] - \\ & \text{sum}[t[i,k,a,c]*v[k,b,i,c],\{k,c\}] + \text{sum}[t[i,c]*t[j,d]*t[k,a]*t[i,b]*v[k,l,c,d],\{k,l,c,d\}] - 2*\text{sum}[t[k,b]*t[i,d]*t[i,j,a,c]*v[k,l,c,d],\{k,l,c,d\}] - \\ & 2*\text{sum}[t[k,a]*t[i,d]*t[i,j,c,b]*v[k,l,c,d],\{k,l,c,d\}] + \text{sum}[t[k,a]*t[i,b]*t[i,j,c,d]*v[k,l,c,d],\{k,l,c,d\}] - 2*\text{sum}[t[j,c]*t[i,d]*t[i,k,a,b]*v[k,l,c,d],\{k,l,c,d\}] - \\ & 2*\text{sum}[t[j,d]*t[i,b]*t[i,k,a,c]*v[k,l,c,d],\{k,l,c,d\}] + \text{sum}[t[j,d]*t[i,b]*t[i,k,c,a]*v[k,l,c,d],\{k,l,c,d\}] - 2*\text{sum}[t[i,c]*t[i,d]*t[j,k,b,a]*v[k,l,c,d],\{k,l,c,d\}] \\ & + \text{sum}[t[i,c]*t[i,a]*t[j,k,b,d]*v[k,l,c,d],\{k,l,c,d\}] + \text{sum}[t[i,c]*t[i,b]*t[j,k,d,a]*v[k,l,c,d],\{k,l,c,d\}] + \text{sum}[t[i,k,c,d]*t[j,l,b,a]*v[k,l,c,d],\{k,l,c,d\}] \\ & + 4*\text{sum}[t[i,k,a,c]*t[j,l,b,d]*v[k,l,c,d],\{k,l,c,d\}] - 2*\text{sum}[t[i,k,c,a]*t[j,l,b,d]*v[k,l,c,d],\{k,l,c,d\}] - 2*\text{sum}[t[i,k,a,b]*t[j,l,c,d]*v[k,l,c,d],\{k,l,c,d\}] - \\ & 2*\text{sum}[t[i,k,a,c]*t[j,l,d,b]*v[k,l,c,d],\{k,l,c,d\}] + \text{sum}[t[i,k,c,a]*t[j,l,d,b]*v[k,l,c,d],\{k,l,c,d\}] + \text{sum}[t[i,c]*t[j,d]*t[k,l,a,b]*v[k,l,c,d],\{k,l,c,d\}] \\ & + \text{sum}[t[i,j,c,d]*t[k,l,a,b]*v[k,l,c,d],\{k,l,c,d\}] - 2*\text{sum}[t[i,j,c,b]*t[k,l,a,d]*v[k,l,c,d],\{k,l,c,d\}] - 2*\text{sum}[t[i,j,a,c]*t[k,l,b,d]*v[k,l,c,d],\{k,l,c,d\}] \\ & + \text{sum}[t[j,c]*t[k,b]*t[i,a]*v[k,l,c,i],\{k,l,c\}] + \text{sum}[t[i,c]*t[j,k,b,a]*v[k,l,c,i],\{k,l,c\}] - 2*\text{sum}[t[i,a]*t[j,k,b,c]*v[k,l,c,i],\{k,l,c\}] \\ & + \text{sum}[t[i,a]*t[j,k,c,b]*v[k,l,c,i],\{k,l,c\}] - 2*\text{sum}[t[k,c]*t[j,l,b,a]*v[k,l,c,i],\{k,l,c\}] + \text{sum}[t[k,a]*t[j,l,b,c]*v[k,l,c,i],\{k,l,c\}] \\ & + \text{sum}[t[k,b]*t[j,l,c,a]*v[k,l,c,i],\{k,l,c\}] + \text{sum}[t[j,c]*t[i,l,k,a,b]*v[k,l,c,i],\{k,l,c\}] + \text{sum}[t[i,c]*t[k,a]*t[i,b]*v[k,l,c,j],\{k,l,c\}] \\ & + \text{sum}[t[i,c]*t[i,k,a,b]*v[k,l,c,j],\{k,l,c\}] - 2*\text{sum}[t[i,b]*t[i,k,a,c]*v[k,l,c,j],\{k,l,c\}] + \text{sum}[t[i,b]*t[i,k,c,a]*v[k,l,c,j],\{k,l,c\}] \\ & + \text{sum}[t[i,c]*t[k,l,a,b]*v[k,l,c,j],\{k,l,c\}] + \text{sum}[t[j,c]*t[i,d]*t[i,k,a,b]*v[k,l,d,c],\{k,l,c,d\}] + \text{sum}[t[j,d]*t[i,b]*t[i,k,a,c]*v[k,l,d,c],\{k,l,c,d\}] \\ & + \text{sum}[t[j,d]*t[i,a]*t[i,k,c,b]*v[k,l,d,c],\{k,l,c,d\}] - 2*\text{sum}[t[i,k,c,d]*t[j,l,b,a]*v[k,l,d,c],\{k,l,c,d\}] - 2*\text{sum}[t[i,k,a,c]*t[j,l,b,d]*v[k,l,d,c],\{k,l,c,d\}] \\ & + \text{sum}[t[i,k,c,a]*t[j,l,b,d]*v[k,l,d,c],\{k,l,c,d\}] + \text{sum}[t[i,k,a,b]*t[j,l,c,d]*v[k,l,d,c],\{k,l,c,d\}] + \text{sum}[t[i,k,c,b]*t[j,l,d,a]*v[k,l,d,c],\{k,l,c,d\}] \\ & + \text{sum}[t[i,k,a,c]*t[j,l,d,b]*v[k,l,d,c],\{k,l,c,d\}] + \text{sum}[t[k,a]*t[i,b]*v[k,l,i,j],\{k,l\}] + \text{sum}[t[k,l,a,b]*v[k,l,i,j],\{k,l\}] \\ & + \text{sum}[t[k,b]*t[i,d]*t[i,j,a,c]*v[l,k,c,d],\{k,l,c,d\}] + \text{sum}[t[k,a]*t[i,d]*t[i,j,c,b]*v[l,k,c,d],\{k,l,c,d\}] + \text{sum}[t[i,c]*t[i,d]*t[j,k,b,a]*v[l,k,c,d],\{k,l,c,d\}] - \\ & 2*\text{sum}[t[i,c]*t[i,a]*t[j,k,b,d]*v[l,k,c,d],\{k,l,c,d\}] + \text{sum}[t[i,c]*t[i,a]*t[j,k,d,b]*v[l,k,c,d],\{k,l,c,d\}] + \text{sum}[t[i,j,c,b]*t[k,l,a,d]*v[l,k,c,d],\{k,l,c,d\}] \\ & + \text{sum}[t[i,j,a,c]*t[k,l,b,d]*v[l,k,c,d],\{k,l,c,d\}] - 2*\text{sum}[t[i,c]*t[i,k,a,b]*v[l,k,c,j],\{k,l,c\}] + \text{sum}[t[i,b]*t[i,k,a,c]*v[l,k,c,j],\{k,l,c\}] \\ & + \text{sum}[t[i,a]*t[i,k,c,b]*v[l,k,c,j],\{k,l,c\}] + v[a,b,i,j] \end{aligned}$$

Simple tensor contraction expression requires

> 13,000 lines of code to implement *sequentially!*

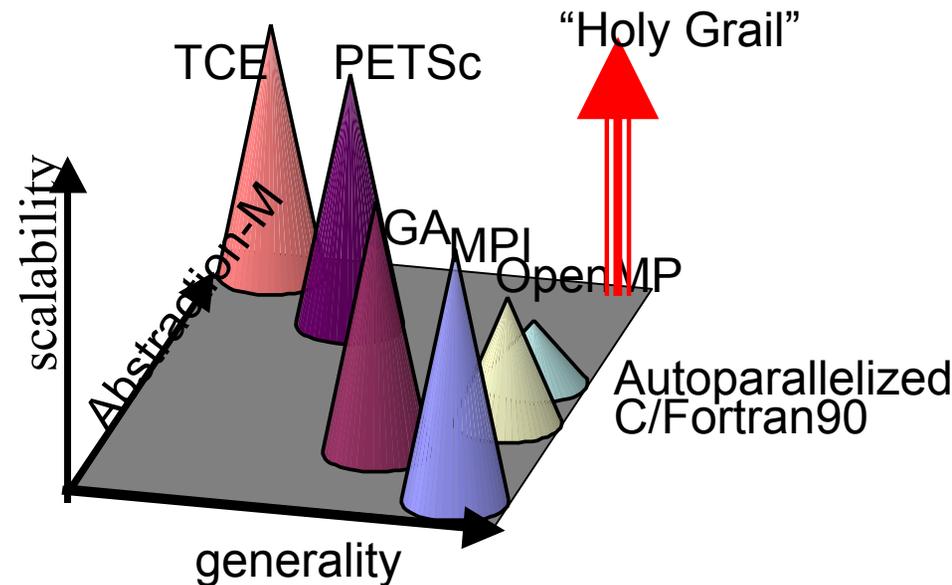
<http://www.cis.ohio-state.edu/~gb/TCE/>

The State of Programming Models

- Programming models haven't changed much in the last five years
 - Mike S
- Programming models haven't changed much in the last *twenty-five* years!
 - (Two-sided) message passing is the “assembly language” of HPC
- Threads and global address space (GAS) models slowing increasing acceptance
 - Global Arrays first developed > 10 years ago
- Applications fighting (vectorizing) compilers
 - Trey W
- How many existing programming model+application provide portable performance from O(4,000 CPUs) to O(40,000 CPUs) for the same problem?
 - CPU counts for 100 TFLOPS

Users Need a Better Way to Write Software for Ultrascale Machines

- The OS should be as transparent as possible
 - Barney M
- User needs “help” *somewhere* in the software stack
- Many software design decisions can be automated *if the software stack knows enough about the problem and the hardware*
- Need to raise level of abstraction at which the user operates



Ten Looking Ahead ~~Five~~ Years

If We Pay Attention to the Programming Environment...

- Ultrascale computers (and beyond) are being used effectively and users love them

If We Continue As Usual...

- Ultrascale computers will be the domain of a few hero programmers
- Most users will prefer their local Beowulf cluster of Windows machines
- Juan M will *still* be complaining about how hard it is to move users up to 512 or more processors