SAND2007-3153P

Overview of the Massively Parallel Quantum Chemistry Program

Core Developers

Sandia National Laboratories

Curtis Janssen
Ida Nielsen
Joe Kenny
Matt Leininger

Virginia Tech

Edward Valeev

Sandia National Laboratories

Outline

- Scalable Computing R & D Dept. Overview
  - Cluster Research: Networks and Storage
  - Performance Tools
  - Software Architecture
  - Applications rethinking the way scientific software is written:
    - Sundance and symbolic PDE solution
    - Particle Simulation Toolkit and Multi-scale modeling
- MPQC
  - Background
  - Performance for DFT, MP2, and LMP2
  - Common Component Architecture work in MPQC

8961 broadly integrates all aspects of high performance computing

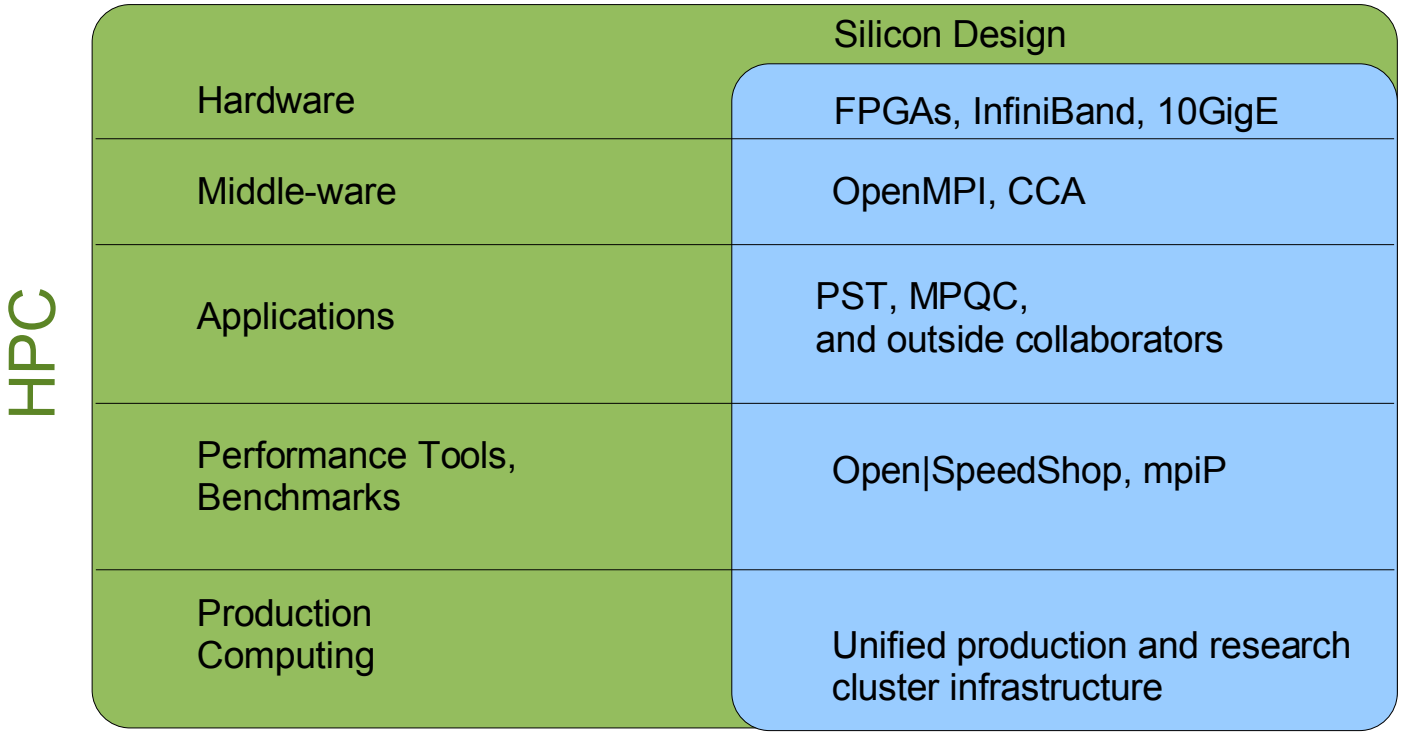

Sandia National Laboratories

Next generation clusters

- Modeling and simulation is a pillar of modern science
  - Replaces experiments and improves our understanding
  - Reliable, high performance, and inexpensive computing is needed to support modeling and simulation
- Sandia has lead the effort to develop the InfiniBand interconnect for use in high performance computing
  - Collaboration between industry, national labs, and universities
  - Sandia leds NNSA ASC PathForward funding for OpenFabrics
  - Started workshop that has grown from 30 to 220 attendees over 4 years
  - OpenFabrics software stack is now in production on many tri-Lab clusters: 12.5K nodes (2Q07) and predict 17K nodes (2Q08)
  - Sandia expertise is sought by outside customers: Wall Street, Oil & Gas, Pharma, etc.
- Remaining challenges include:
  - Provide foundations for applications scaling up in
    - number of nodes
    - number of cores per node
  - Load balancing of network traffic/routing
  - Better support for scientific computing
  - Virtual machines and I/O

OPENFABRICS ALLIANCE

Sandia National Laboratories

Advanced networking for storage and distance computing

- Large scale applications can generate hundreds of TeraBytes of data
  - Data must be saved in persistent storage
  - Data must be moved to local sites for interpretation and visualization
  - Storage and wide-area networks are limiting factors in large scale computation
- A variety of approaches can be used to address storage
  - Hardware accelerated protocol processing reduces the processing requirements for sending and receiving data
  - Remote direct memory access eliminates operating system intervention and data copying during I/O operations
- Future work:
  - Investigate various 10Gbps and beyond networking technologies in network storage
  - Develop evaluation and extrapolation methodologies to give confidence in modeling and simulation of storage for future platform designs
  - Apply virtual machine technology to both storage and Storage Intensive Super Computing

Sandia National Laboratories

Performance and programming tools

- Both architectures and software are changing, and developers need to understand performance to get good utilization
  - Architectures shifting from faster processors to more cores and more complex memory hierarchies
  - Software is becoming more sophisticated and complex
    - Understanding performance and debugging is more difficult
  - A (very modest) 20% performance improvement on a $20M machine is effectively a savings of $5M in purchase (plus maintenance savings)
- 8961 is working with LLNL and LANL on next generation tools
  - Performance Measurement: Open|SpeedShop, mpiP
  - Debugging: Valgrind (memory and thread debugging)
- Future avenues of work:
  - Performance modeling and prediction
    - supporting hardware architecture design
  - Tools that assist in the interpretation of performance data
  - Programming models for next generation machines and applications

Software architectures

- Effective software engineering is the key to world-class computational science
  - Complex systems are modeled by complex code
    - Components are the accepted means to manage complexity
    - Common Component Architecture is dedicated to high performance computing
  - Full realization of scientific potential requires innovative simulation capability
  - Complete solutions unlikely from “hero programmers” or even “hero groups”
  - Progress requires large-scale collaborations within and between domains
- 8961 is providing leadership in both development and application of Common Component Architecture (CCA) work to address these challenges through two SciDAC2 projects.
  - CCA core development
  - CCA components for quantum chemistry
- Future work:
  - Improved solver components, components for interfacing quantum mechanics and molecule mechanics

Overview of Sundance/Nihilo

- Sundance is a package of high-level components for building parallel PDE simulations
- User builds problem description using high-level symbolic objects
  - Just write the equation set!
- Automatic functional differentiation built in → optimization & sensitivity analysis
- Abstract interfaces for atomistic-continuum coupling
  - In use for multiscale simulation of nanofluidic ion transport in a pore
- Easy to write a high-performance multiphysics simulator

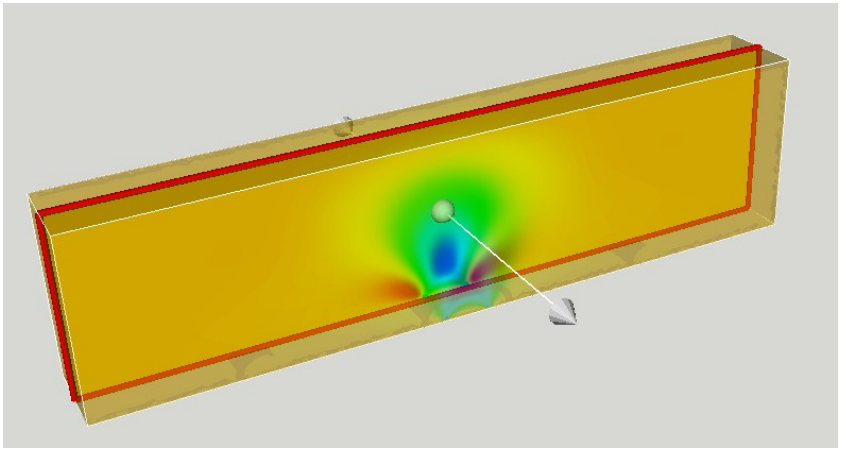

Flow velocity above an electrode gap in a microfluidic channel. Sundance was used to develop a simulator coupling electrical, thermal, and fluid effects. Time between receipt of the paper describing the problem's physics to the creation of a working simulation code was under 24 hours
(KL and B. van Bloemen Waanders, 2006)

Sandia National Laboratories

Particle Simulation Toolkit (PST) and multi-scale modeling

- Enormous improvements in computational power have dramatically increased the scope of what is computable
  - Traditional scientific computing software fails in:
    - does not efficiently address the large range of scales
    - handling coupled nature of current problems
- We have developed highly effective, problem-oriented multi-scale computing methodologies
  - Example application: nanopore studies of separations, transport efficiency and selectivity

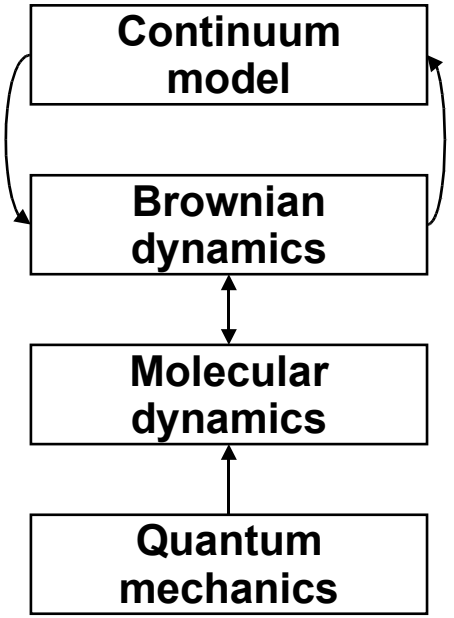

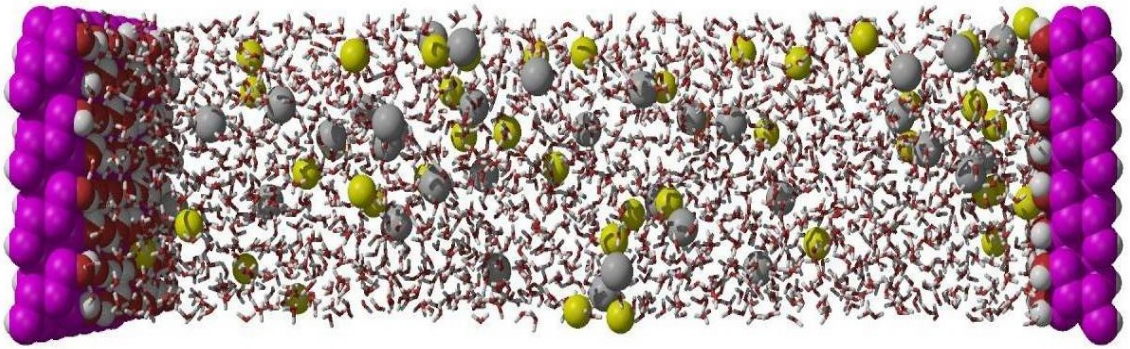

Sandia National Laboratories

The Massively Parallel Quantum Chemistry program (MPQC)

- MPQC is a scalable, extensible QC package
  - Standard methods such as Hartree-Fock and Density-Functional Theory
  - Highly scalable perturbation theory and explicitly correlated methods
  - Parallelized to permit large-scale calculations
  - Local correlation methods to extend range of applicability
  - Development of chemistry CCA components
- Hybrid multi-threading/distributed memory testbed
- Implementation
  - 370,000 lines of C++
  - Object-oriented design
  - Parallelized from the outset

Sandia National Laboratories

DFT Performance

- B3LYP hybrid functional
- Uracil Dimer
- 264 basis functions
- Dual 3.06GHz Intel Xeon nodes
- Infiniband 4X interconnect
- Total time on 70 nodes:
  - 38.74 seconds

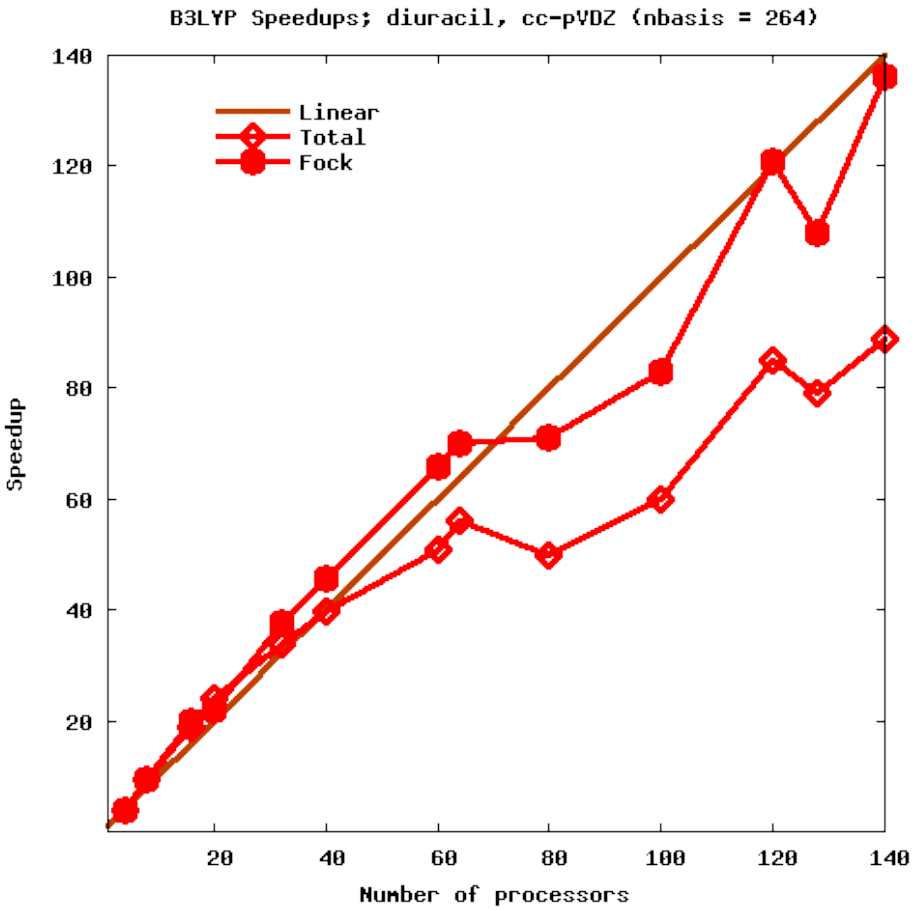

Sandia National Laboratories

MP2 Performance

- Uracil Dimer
- 264 basis functions
- Dual 3.06GHz Intel Xeon nodes
- Infiniband 4X interconnect
- Total time on 70 nodes:
  - 101 seconds

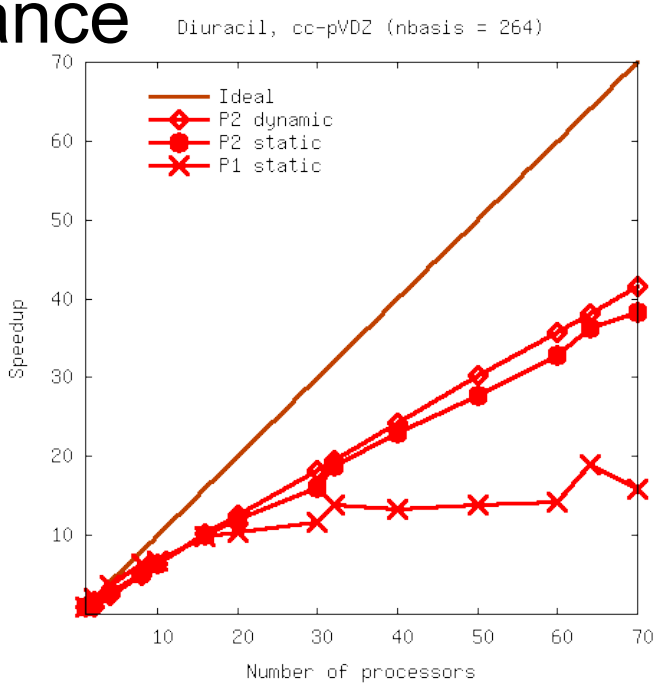

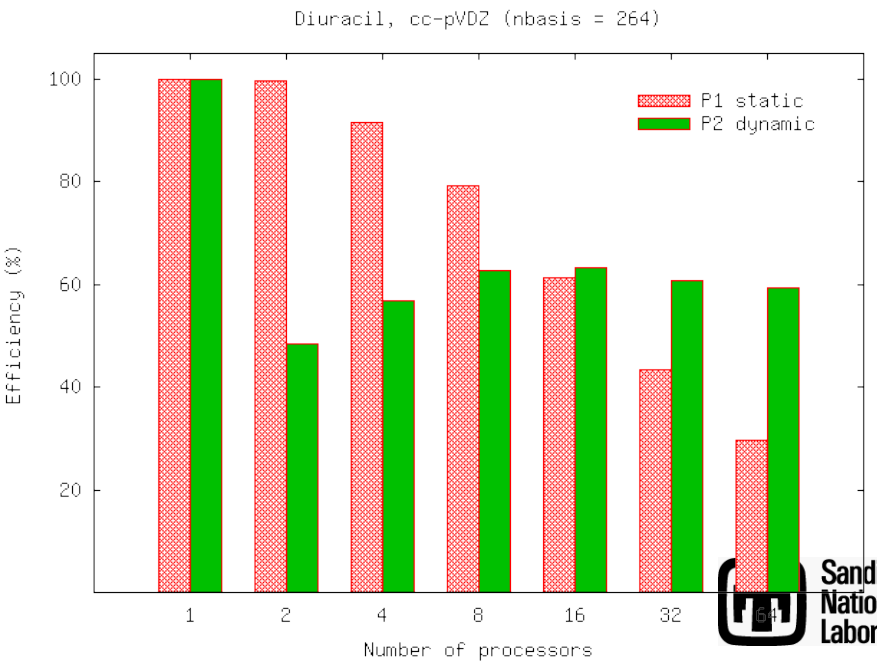

LMP2 Scaling and Performance

- System scaling runs: CnH2n+2
- Parallel scaling runs: C32H66
- cc-pVDZ basis
- Dual 3.6 Ghz Intel Xeon nodes
- Infiniband 4X

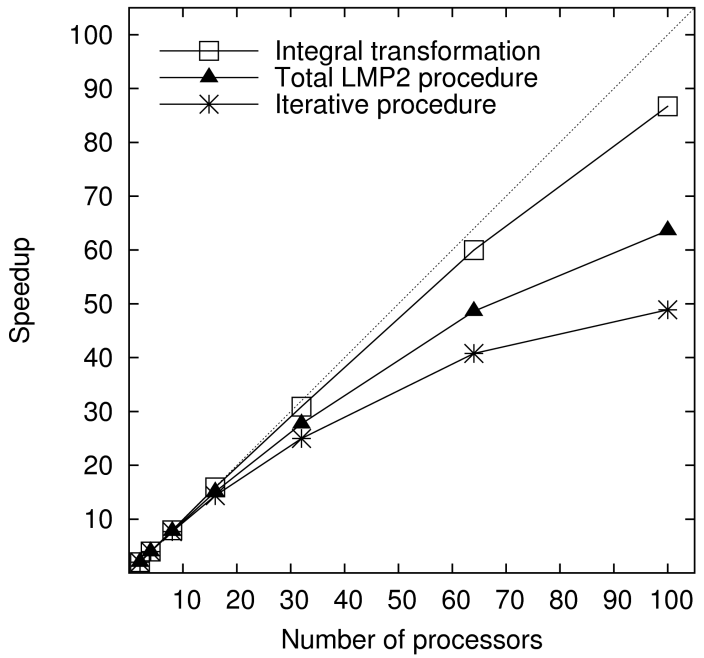

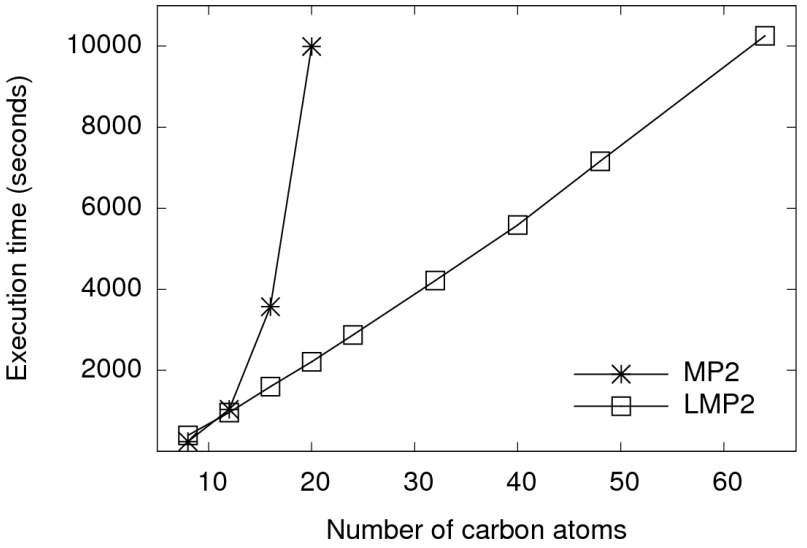

Sandia National Laboratories

Component Architectures for Quantum Chemistry: Forging New Capabilities and Insights

CCA/Chemistry team members and collaborators

Sandia National Laboratories

Joe Kenny
Curtis Janssen
Ida Nielsen
Rob Armstrong
Ben Allan

Pacific Northwest National Laboratory
Operated by Battelle for the U.S. Department of Energy

Yuri Alexeev
Manojkumar Krishnan
Elizabeth Jurrus
Carl Fahlstrom
Jarek Nieplocha

NWChem
high performance computational chemistry software

Gary Kumfert

Steve Benson
Jason Sarich
Lois Curfman McInnes

ornl
OAK RIDGE NATIONAL LABORATORY

David Bernholdt
Ricky Kendall

AMES LABORATORY

GAMESS

Theresa Windus
Sasha Sosonkina
Fang Peng
Mark Gordon

VirginiaTech
Invent the Future

Edward Valeev

Sandia National Laboratories

State of modeling & simulation software

Simulation sciences are maturing

- Mainstream research tool
- Reliable and on par with experiment

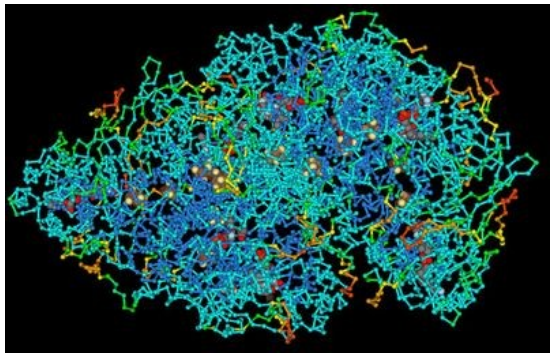

Many fundamental problems remain intractable

- Accurate simulations at biologically interesting scales
- Progress requires innovative, multi-disciplinary, multi-scale approaches

- Full realization of scientific potential requires innovative simulation capability
- Complete solutions unlikelyfrom “hero programmers” or even “hero groups”
- Progress requires large-scale collaborations within and between domains

Effective software engineering is key to world-class computational science

Sandia National Laboratories

Scientists as programmers

- Not known for sound software engineering practices
  - Simplest approach often preferred, even if inefficient
  - Coding practices are often out-dated
    - Poor style, little documentation, incomplete testing
    - Difficult to convert to modern programming techniques
      - Learning curve, poor training, and legacy code are issues.
  - Common programming tools that helped make large-scale scale software efforts such as GNU/Linux successful are not uniformly utilized:
    - minimal use of software configuration management, build systems
- Diverse community of government/academic, noncommercial quantum chemistry (QC) packages
  - Limit ability to leverage existing capabilities
  - Interaction between QC and other fields even more difficult

Diverse computer architectures are another complicating factor

- Supporting N advanced computer architectures requires > N times the effort of running only on my laptop
- Currently at another junction in computer architecture: can no longer rely on clock speed increases to improve performance

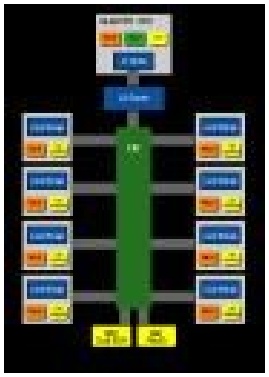

Heterogeneous Multicore “Cell” Processors

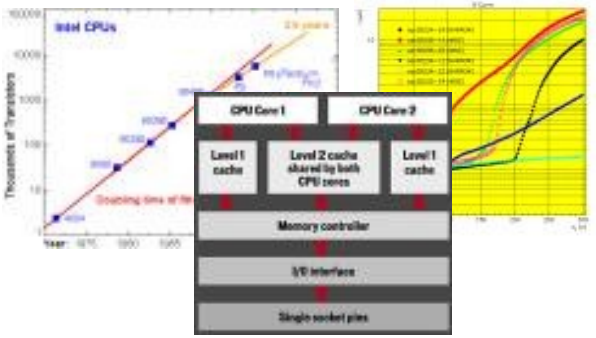

Multicore Processors

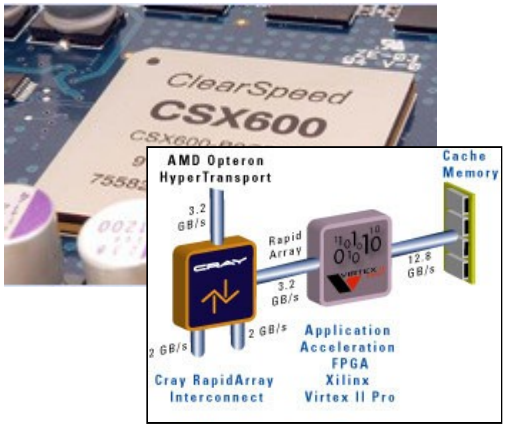

Accelerator Cards, FPGAs

Sandia National Laboratories

Diversity isn't a bad thing

- Diversity in software
  - Allows exploration of alternative solutions
  - Allow individuals with different skill sets to participate
  - Perfect coordination takes unlimited effort
- Diversity in architectures
  - Opportunity to find high levels of performance with less cost and lower power requirements
  - Different algorithms have different architecture needs
    - Hetergeneous solutions may eventually be a part of the solution

But must avoid the monolithic code trap

- Total effort is divided by duplicating capabilities
- Limits the quality of capabilities

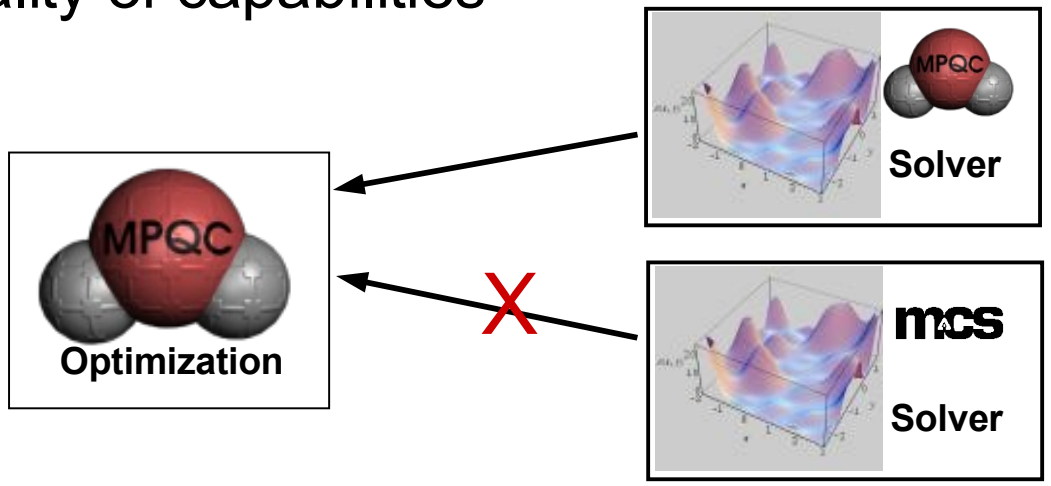

- Limits the capabilities

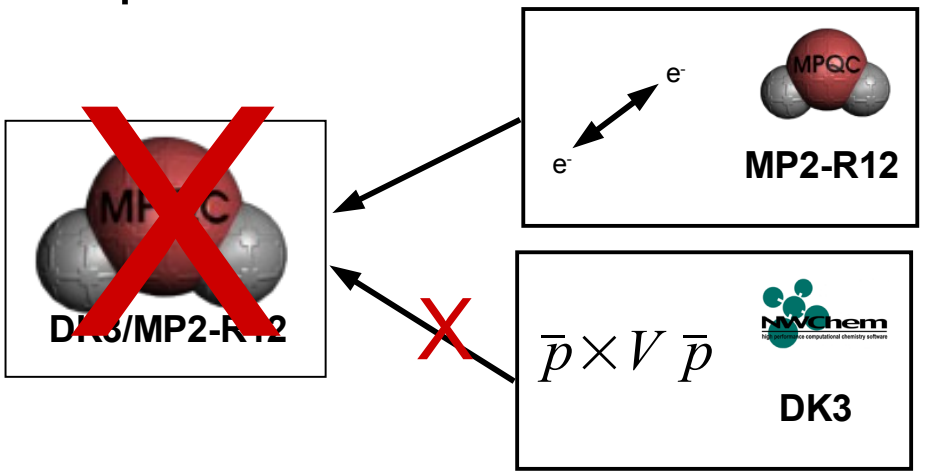

Sandia National Laboratories

How can human effort scale in this diverse software/architecture environment?

- Object-oriented methodologies? There are issues:
  - Cannot leave out legacy codes
  - Even with modern codes, design patterns may be similar but implementation/language is not. Code bases incompatible at a low level.
  - Not a complete solution
- Characteristics of a solution:
  - Must support multiple languages
  - Must allow for mostly independent programming in packages using it
  - Community must agree on a few well-defined or common elements in the design that place minimal constraints on each software package
  - The Common Component Architecture is designed to satisfy these requirements.

Illustration of complications in QC: Chromium hydroxides

- Accurate thermochemical knowledge needed to understand contamination in industrial settings and pollution
- Experimental data is missing or inconsistent
- Six reactions used to obtain heat of formation for $Cr(OH)_n$, n = 2–6 and $CrO(OH)_4$.

$$Cr(OH)_2 + Cr \rightarrow 2\ CrOH$$

$$Cr(OH)_3 + 2\ Cr \rightarrow 3\ CrOH$$

$$Cr(OH)_4 + 3\ Cr \rightarrow 4\ CrOH$$

$$Cr(OH)_5 + CrOH \rightarrow CrOH$$

$$Cr(OH)_6 \rightarrow CrO_3 + 3\ H_2O$$

$$CrO(OH)_4 \rightarrow CrO_3 + 2\ H_2O$$

Sandia National Laboratories

High accuracy is hard

- Thousands of hours of CPU time and four quantum chemistry code suites later ...

	$Cr(OH)_2$ Rxn 1	$Cr(OH)_3$ Rxn 2	$Cr(OH)_4$ Rxn 3	$Cr(OH)_5$ Rxn 4	$Cr(OH)_6$ Rxn 5	$CrO(OH)_4$ Rxn 6
ΔE_{rxn}[HF]	7.71	−4.60	−50.63	132.49	−29.21	24.02
δ[CCSD]	+7.17	+18.22	+40.47	−44.51	+15.57	+1.45
δ[CCSD(T)]	+0.36	+4.62	+13.06	−11.03	+9.62	+0.78
δ[basis]	−0.30	−1.15	−2.11	−3.49	−3.91	−2.63
δ[core]	−0.37	+0.76	+1.82	[b]	+2.08	+0.90
δ[rel]	−2.53	−2.87	−3.19	+3.89	+2.06	+3.42
δ[ZPVE]	−1.06	−2.64	−4.12	−5.53	−8.07	−4.99
$\Delta H^\circ_{rxn,0}$	10.98	12.34	−4.70	71.82	−11.86	22.95
$\Delta H^\circ_{f,0}$	−72.08	−151.21	−211.97	−241.74	−235.85	−213.55
$\Delta H^\circ_{f,298.15}$	−73.19	−153.36	−214.94	−247.07	−242.64	−218.09

- Limited by abilities of each code
  - Assumed additive contributions for different effects
  - Choice of methods not always optimal

Table from: Nielsen, Allendorf, J. Phys. Chem. A, **110**, p4093, 2006

Sandia National Laboratories

What made this problem so hard?

- Different program suites have different strengths
  - Some overlap, but important differences in supported methods
  - Different numerical properties
  - Different levels of support for various architectures
  - Need better ways of interchanging program suites and sharing capabilities between suites
- Gets even harder when quantum chemistry is a component of multi-scale, multi-physics computations
  - Building applications that rely on multiple application domains is even more complex
  - Need ability to export and import capabilities

Component architectures are designed to address these problems

- Language neutral interface specification
  - Different code teams focus only on the common interface
  - Use SIDL: Scientific Interface definition language
- Provides a runtime environment
  - Can dynamically compose an application

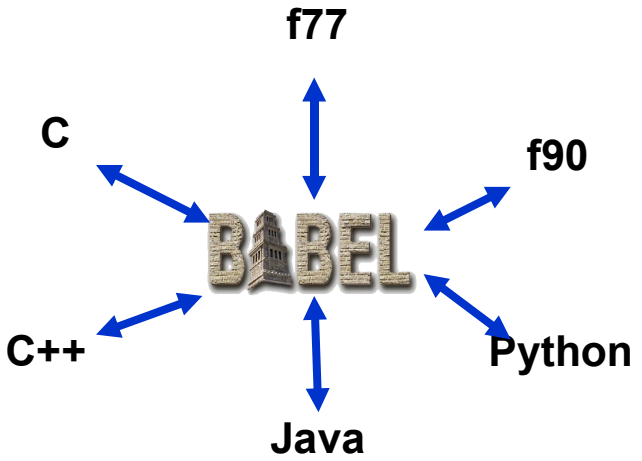

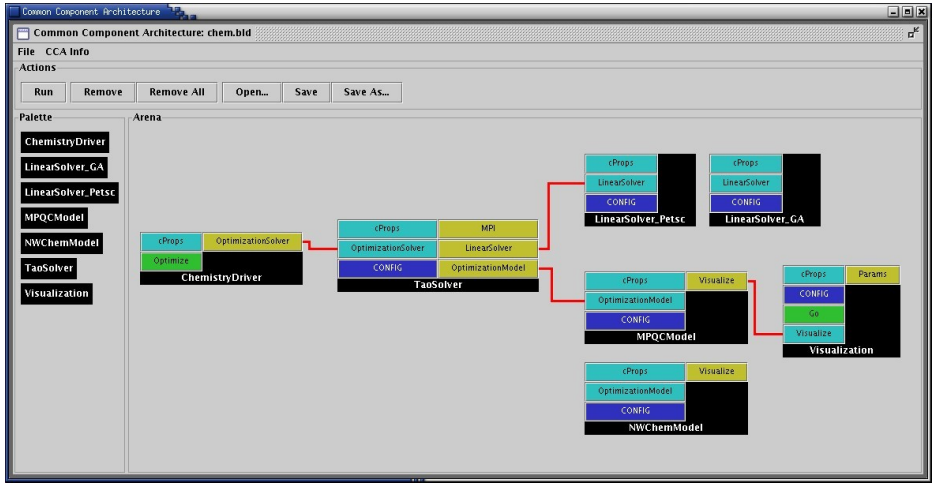

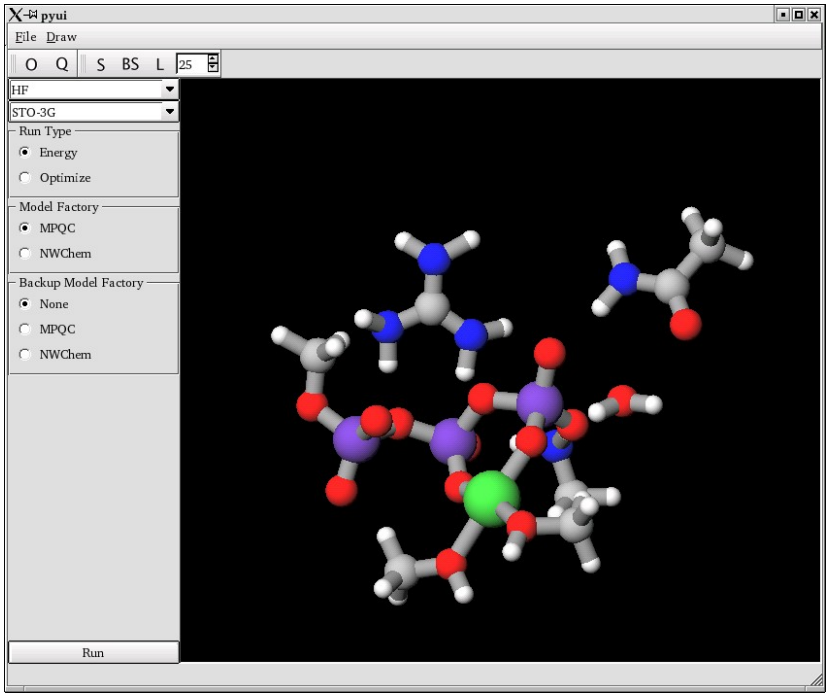

Two applications of the Common Component Architecture

- High-level components for geometry optimization
- Low-level components for enabling new methods

Sandia National Laboratories

High level components and their use in geometry optimization

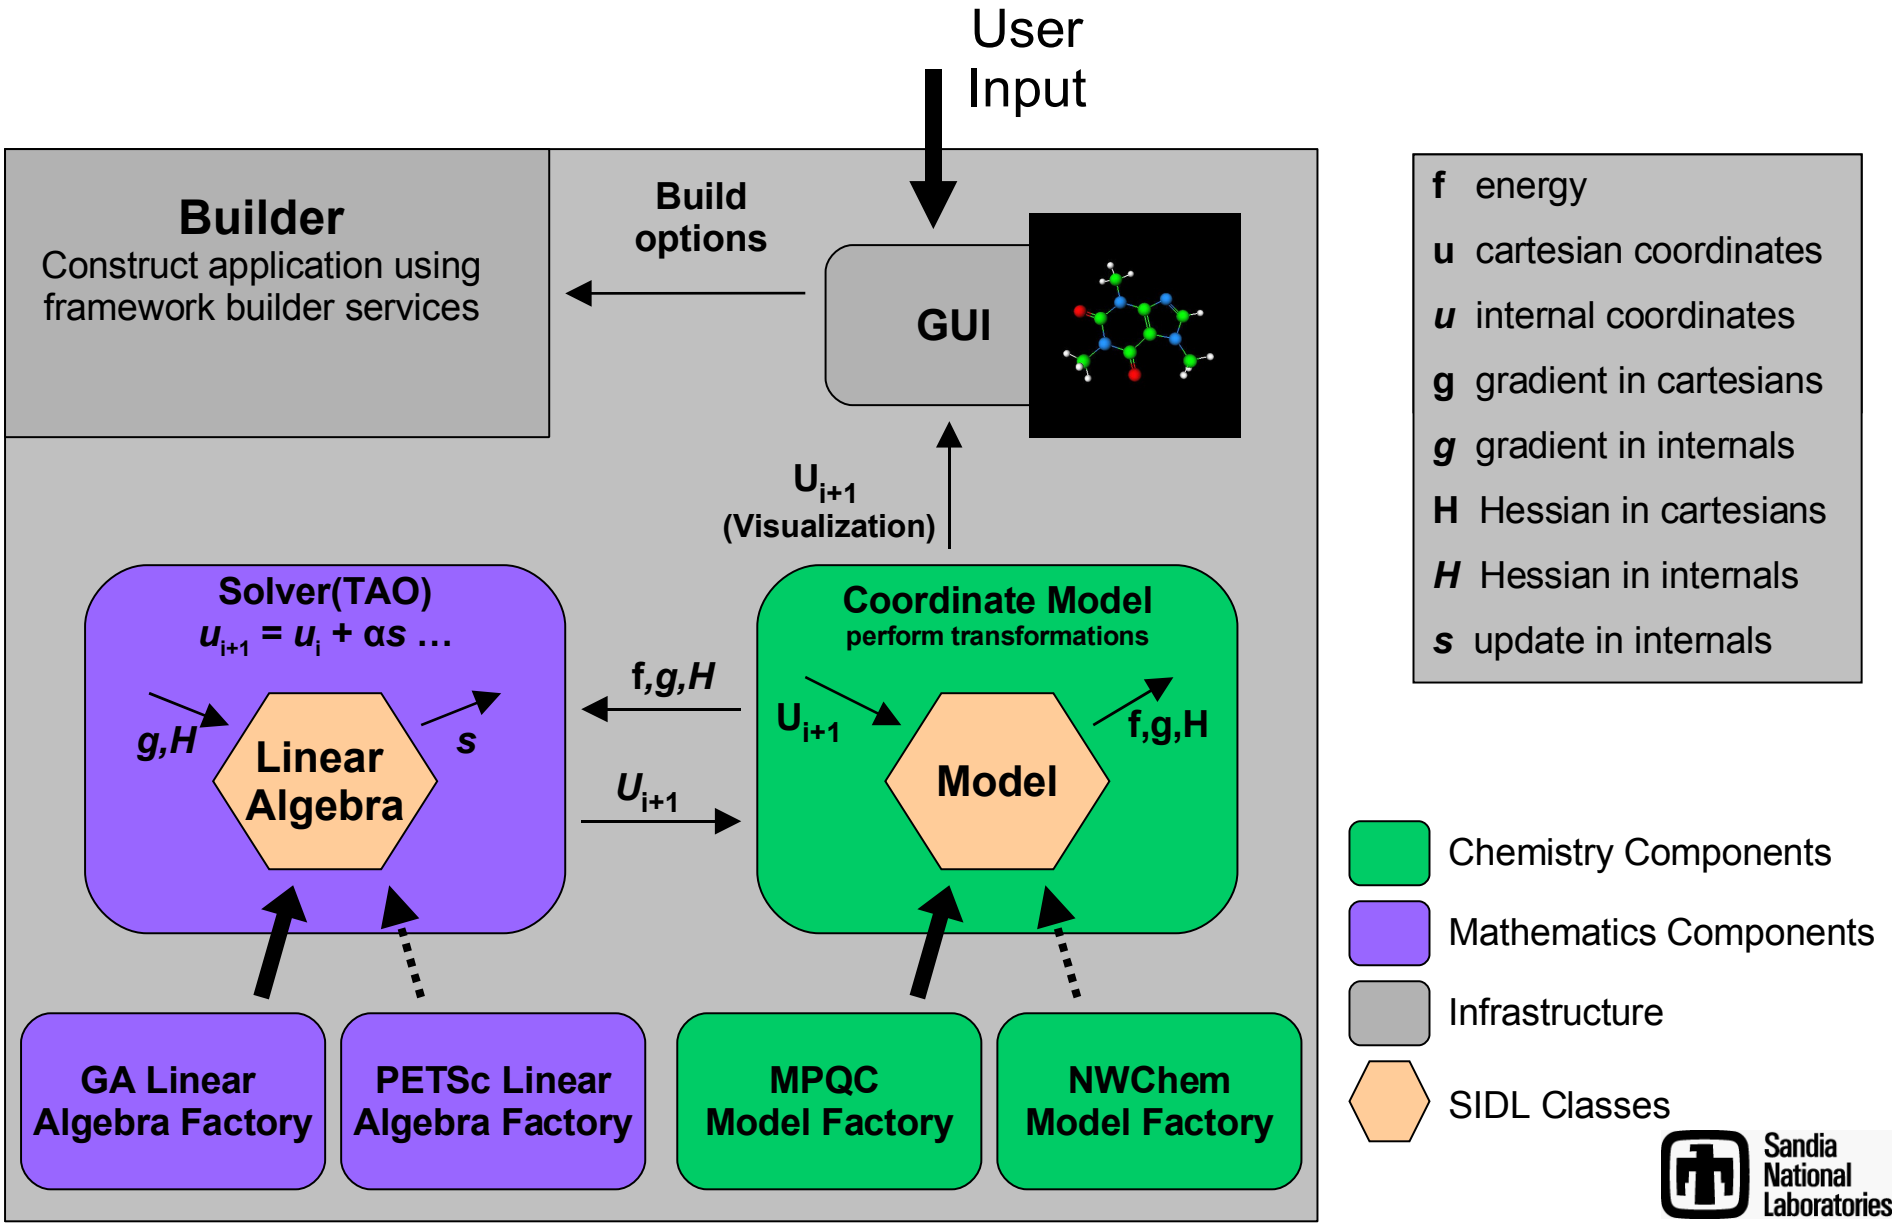

Sandia National Laboratories

Enabled direct comparison of various solvers for molecular structures

QC Package	MPQC	MPQC	NWChem	NWChem	NWChem
Algorithm	BFGS	TAO/LMVM	BFGS	BFGS	TAO/LMVM
Line Search	no	yes	no	yes	yes
Guess Hessian	unit	scaled unit	0.5*unit	0.5*unit	scaled unit
Glycine ($C_2H_5NO_2$)	26/26	19/19 +27%	33/33	65/33	19/19
Isoprene (C_5H_{10})	75/75	43/43 +43%	56/56	89/45	45/45
Phosphoserine ($C_3H_8NO_6P$)	85/85	62/62 +27%	79/79	121/61	67/67
Acetylsalicylic Acid ($C_9H_8O_4$)	54/54	48/48 +21%	43/43	83/42	51/51
Cholesterol ($C_{27}H_{46}O$)	27/27	30/30 -11%	33/33	—/—	30/30

Stand-alone MPQC/NWChem

TAO Solver Component

Number of energy/gradient evaluations required to determine minimum energy structure

Integration gave us insights into problems with our solvers ... and a new solver

Sandia National Laboratories

Low-level components to extend capabilities of programs

- Integrals of many operators are at the core of quantum chemistry programs:

$$\int dr_1 dr_2 \phi_1(r_1)\phi_2(r_1) r_{12} \phi_3(r_2)\phi_4(r_2)$$

$$\int dr\, \phi_1(r)\nabla^2\phi_2(r)$$

$$\phi_i(r) = x_i^a y_i^b z_i^c e^{-\alpha_i (r-R_i)^2}$$

$$\int dr_1 dr_2 \phi_1(r_1)\phi_2(r_1)[\nabla_1^2, r_{12}]\phi_3(r_2)\phi_4(r_2)$$

$$\int dr_1 dr_2 \phi_1(r_1)\phi_2(r_1)\frac{1}{r_{12}}\phi_3(r_2)\phi_4(r_2)$$

- Integrals programs do not implement all integral types
- Ability to share integrals and combine packages will
  - enable new science
  - permit selection of most efficient package for each machine

Low-level components provide an extreme test of the CCA

- Low-level components tend to be finer grained with more function call overhead. For Hartree-Fock:

		MPQC	CCA	Overhead
H_2O cc-pVQZ	energy	19.9	21.0	5.5%
	gradient	39.5	41.3	4.6%
C_5H_{10} cc-pVDZ	energy	86.3	93.7	8.6%
	gradient	204.8	219.1	7.0%

This is a worst case scenario:

- More sophisticated methods, vector interfaces, and non-direct methods = less overhead
- Latest version of the CCA software improves results

Sandia National Laboratories

Using the integral components to develop a new method

- Douglas-Kroll allows simple relativistic effect inclusion:

$$h_1^{sf} = c\left(p^2 + c^2\right)^{1/2} - c^2 + A_p V A_p + B_p \bar{p}\cdot V\, \bar{p}\, B_p + \ldots + F_p \bar{p}\times V\, \bar{p}\, Y_p \bar{p}\times V\, \bar{p}\, F_p + \ldots$$

special integral types

- r_{12} methods allow more rapid wfn convergence

$$\Psi^{(1)}_{\text{MP2-R12}} = d^{ij}_{ab} a^{ab}_{ij} \Phi + c^{ij}_{\boldsymbol{kl}} \bar{R}^{kl}_{\alpha\beta} a^{\alpha\beta}_{ij} \Phi$$

special integral type

Combination of these methods had not been done

- Even though ideal for high Z core correlation

Sandia National Laboratories

Missing piece: component to combine multiple integral packages

Architecture:

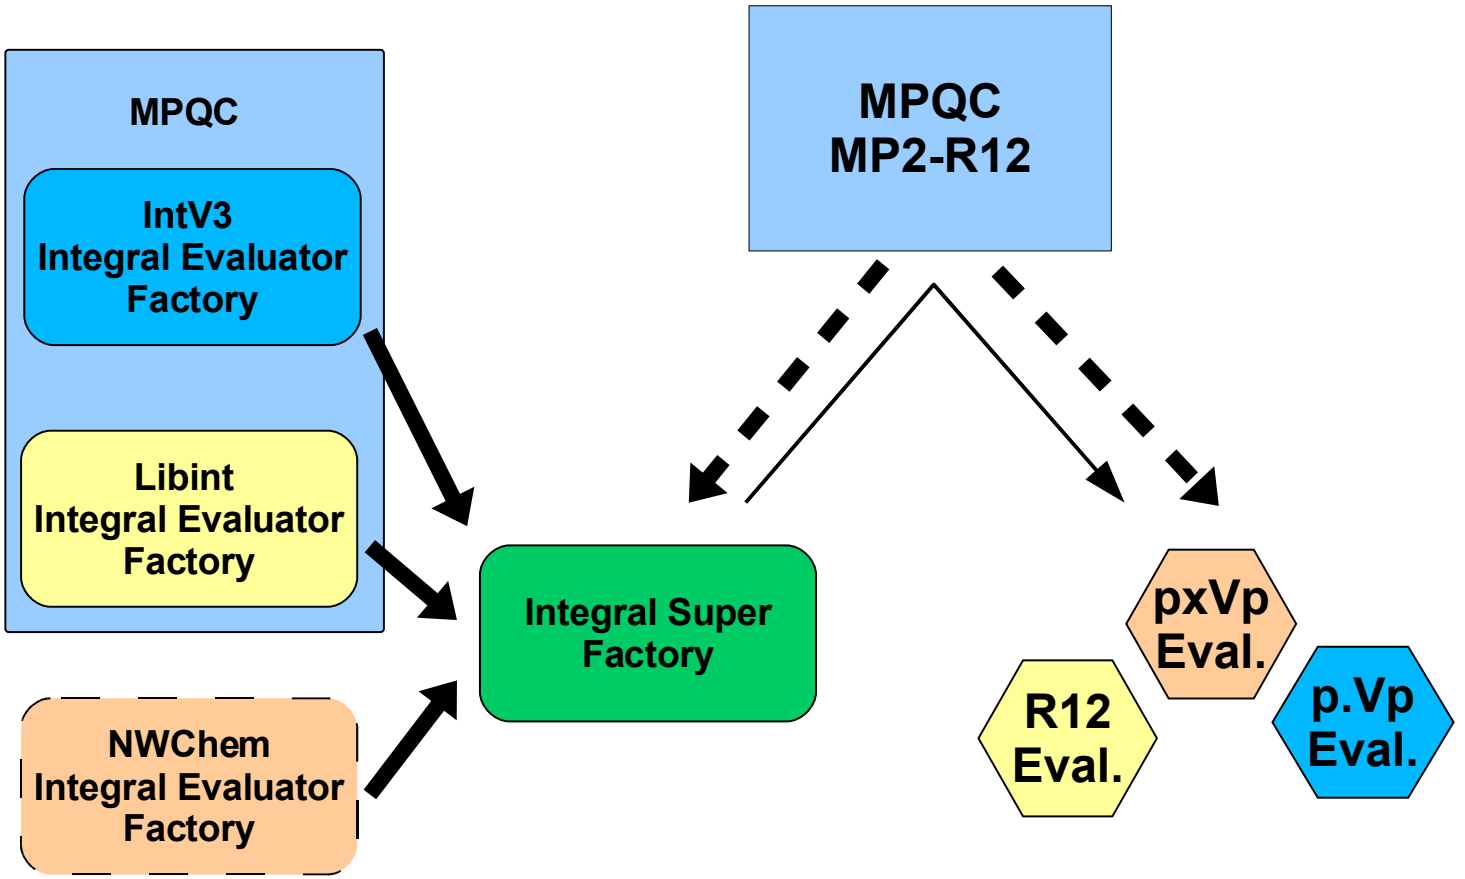

- Example of where significant functionality gets implemented into component specific code

Applying this to the chromium hydroxide example

- Opportunity to combine three corrections in to one:

  δ[core] + δ[rel.] + δ[basis] → δ[core+rel.+basis]

	$Cr(OH)_6$ Reaction 1		$CrO(OH)_4$ Reaction 2	
ΔE_{rxn}[HF]	−29.66		23.58	
δ[MP2]	+9.81		−7.47	
δ[CCSD]	+5.04		+8.45	
δ[CCSD(T)]	+9.94		+0.96	
	MP2	MP2-R12	MP2	MP2-R12
δ[basis]	−3.55		−2.39	
δ[core]	−0.32		−0.79	
δ[rel]	+2.08		+3.44	
δ[basis+core+rel]	**(Σ = -1.79)**	−1.85	**(Σ = 0.33)**	+0.32
δ[ZPVE]	−8.07	−8.07	−4.99	−4.99
$\Delta H^\circ_{rxn,0}$	−14.73	−14.79	20.79	20.85
$\Delta H^\circ_{f,0}$	−232.98	−232.92	−211.40	−211.45
$\Delta H^\circ_{f,298.15}$	−239.77	−239.71	−215.94	−216.00

Sandia National Laboratories

A cautionary note

- Components greatly increase the flexibility of programming
  - Make it easier for non-experts to construct applications
- But must understand whether or not the applications are valid
- Example from this case:
  - MP2-R12 requires matrix elements of:

$$[\hat{f}, r_{12}] \overset{\text{n.r.}}{\rightarrow} [-\nabla^2/2, r_{12}] + [K, r_{12}]$$

$$[\hat{f}, r_{12}] \overset{\text{rel}}{\rightarrow} [c(p^2+c^2)^{1/2} + A_p V A_p + B_p \bar{p}\cdot V \bar{p} B_p + \ldots + K, r_{12}]$$

likewise for $[r_{12}, [\hat{f}, r_{12}]]$

- Must compute additional terms, or quantify error
  - Sophisticated software architectures do not eliminate the need for experts, but they do improve the expert's productivity

Future Work

- Planned:
  - Common Component Architecture work
    - Quantum mechanics/molecular mechanics interface & implementation
    - Effective Fragment Potential
    - General one body operator interfaces (solvation, for example)
- Areas of interest:
  - Fully distributed Fock build for memory-starved machines
  - Reduced scaling HF/DFT
  - High-level correlation with Tensor Contraction Engine
  - Exploration of high accuracy methods for materials (beyond DFT)
  - Properties
  - Manycore performance issues

Sandia National Laboratories

For more information

- MPQC Home Page
  - http://www.mpqc.org
- MPQC Project Page
  - http://www.sf.net/projects/mpqc
    - Bug tracking
    - Mailing lists
    - CVS repository
    - Code releases
- CCA Home Page
  - http://www.cca-forum.org
- Send questions to
  - cljanss@sandia.gov

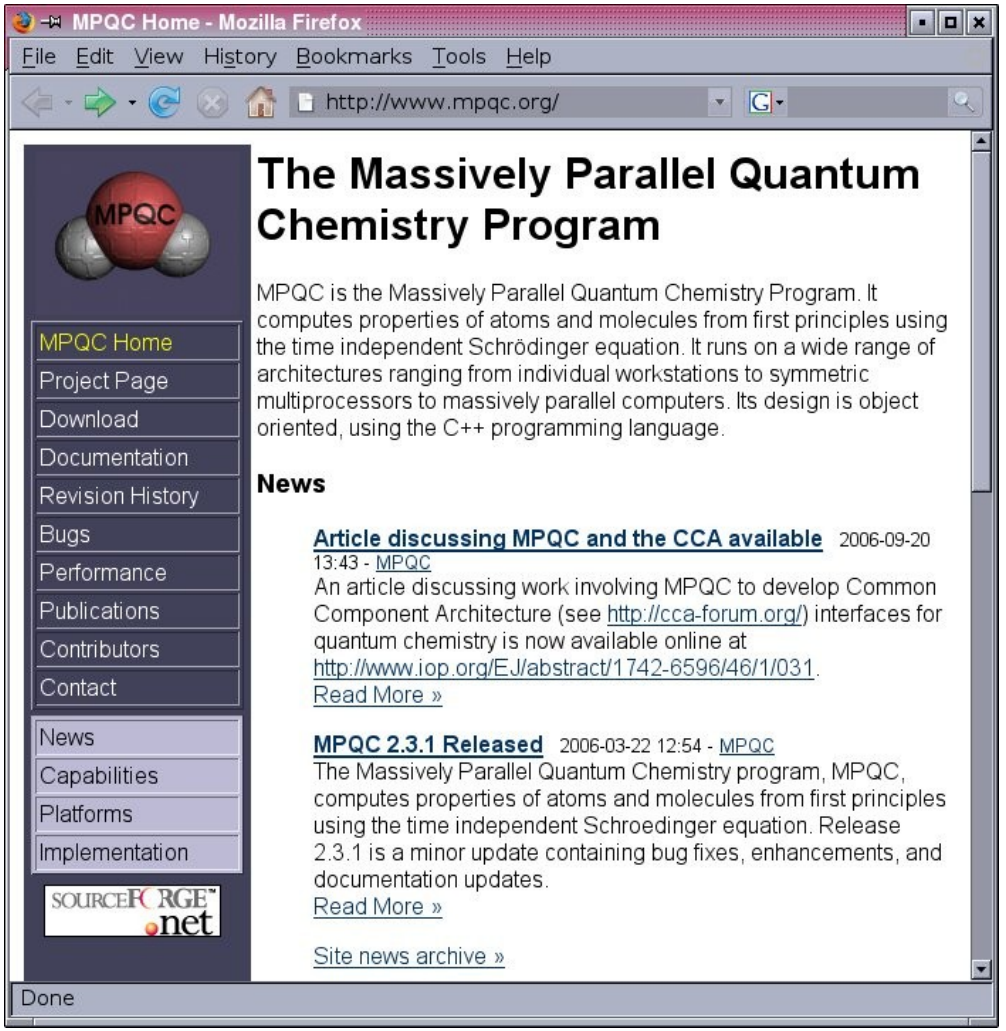

Sandia National Laboratories