Experimental Mathematics, Multicore Processors and Highly Parallel Computing

David H Bailey

Lawrence Berkeley National Laboratory

`http://crd.lbl.gov/~dhbailey`

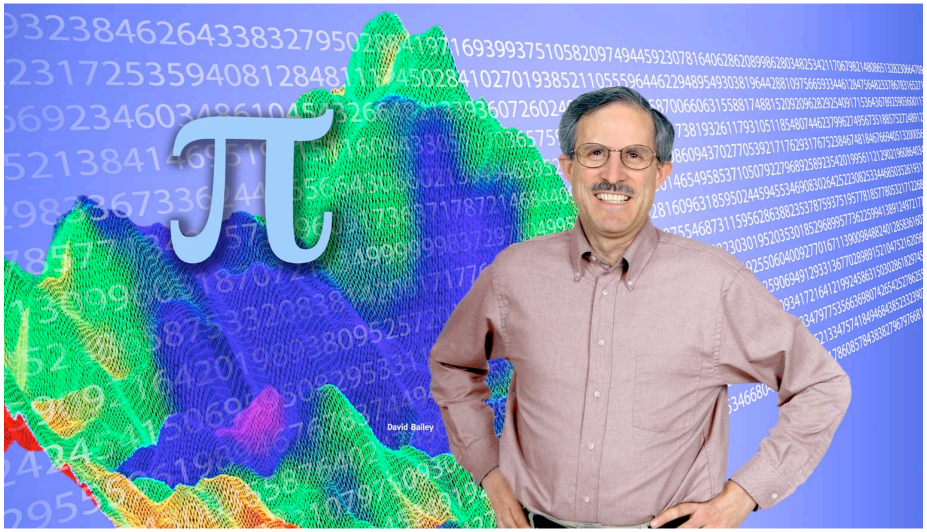

Petascale-Multicore Systems: Which Applications Will Run Well?

Likely to run well on petascale systems with multicore processors:

- Applications with enormous natural concurrency: ~10^8-way concurrency at every significant step of the computation.
- Applications with mostly local data access: inner kernels have substantial computation on relatively little data.
- Applications with two natural levels of parallelism: distributed at high level, shared memory at low level.

Other classes of applications are not likely to run at optimal rates.

Major challenges for the foreseeable future:

- Finding (and exploiting) huge levels of concurrency in applications.
- Minimizing communication, via improved algorithms and software.
- Tuning local-node code for multicore processors.

Sparse Matrix Test Suite for Multicore Tuning

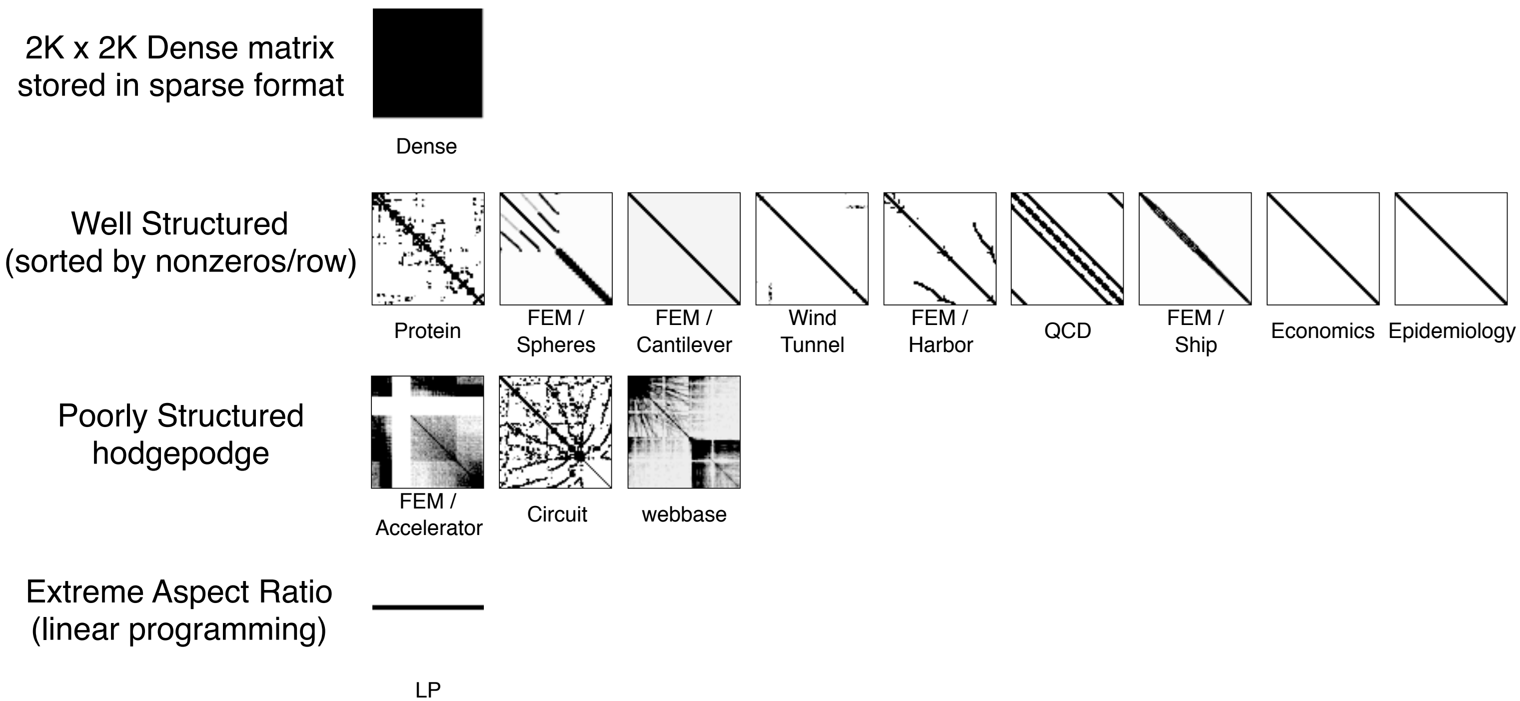

In the next few slides, performance results on this suite will be shown for various multicore systems, using various prototype semi-automatic tuning schemes. For full details, see this paper:

Samuel Williams, Kaushik Datta, Jonathan Carter, Leonid Oliker, John Shalf, Katherine Yelick, DHB, "PERI: Auto-tuning Memory Intensive Kernels for Multicore," *Journal of Physics: Conference Series*, vol. 125 (2008), pg. 012038; available at: http://crd.lbl.gov/~dhbailey/dhbpapers/scidac08_peri.pdf

Four Multi-Core SMP Systems

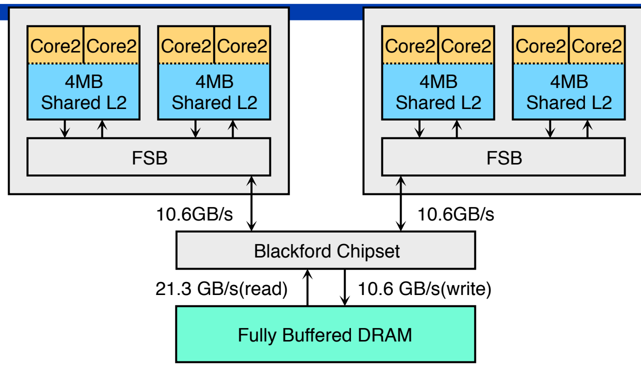

Intel Clovertown

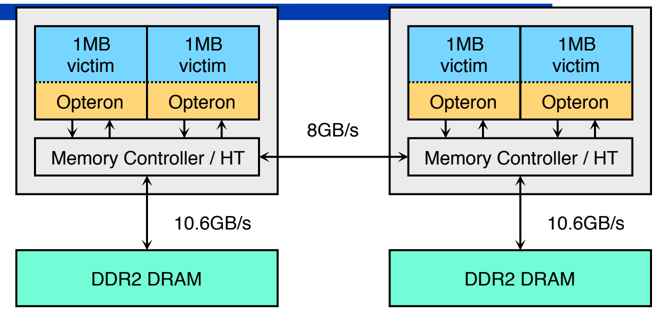

AMD Opteron

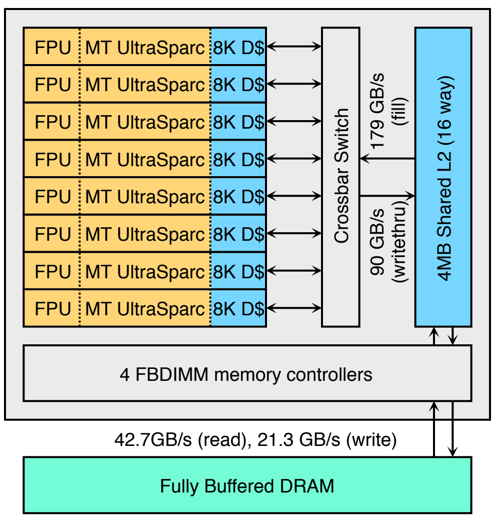

Sun Niagara2

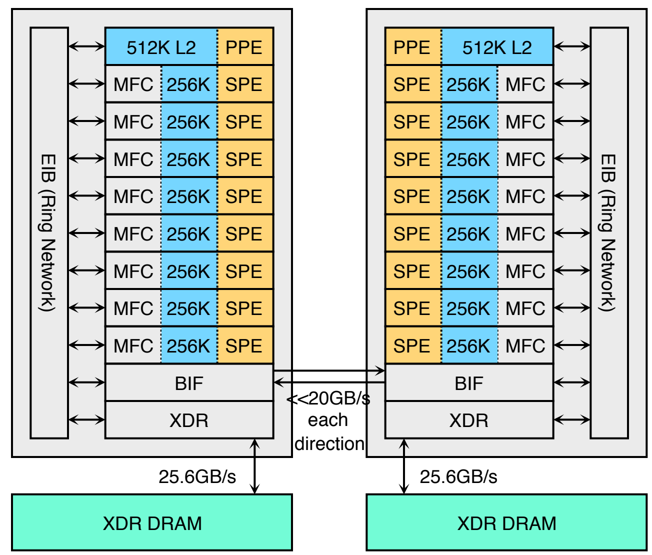

IBM Cell Blade

Four Multi-Core SMP Systems

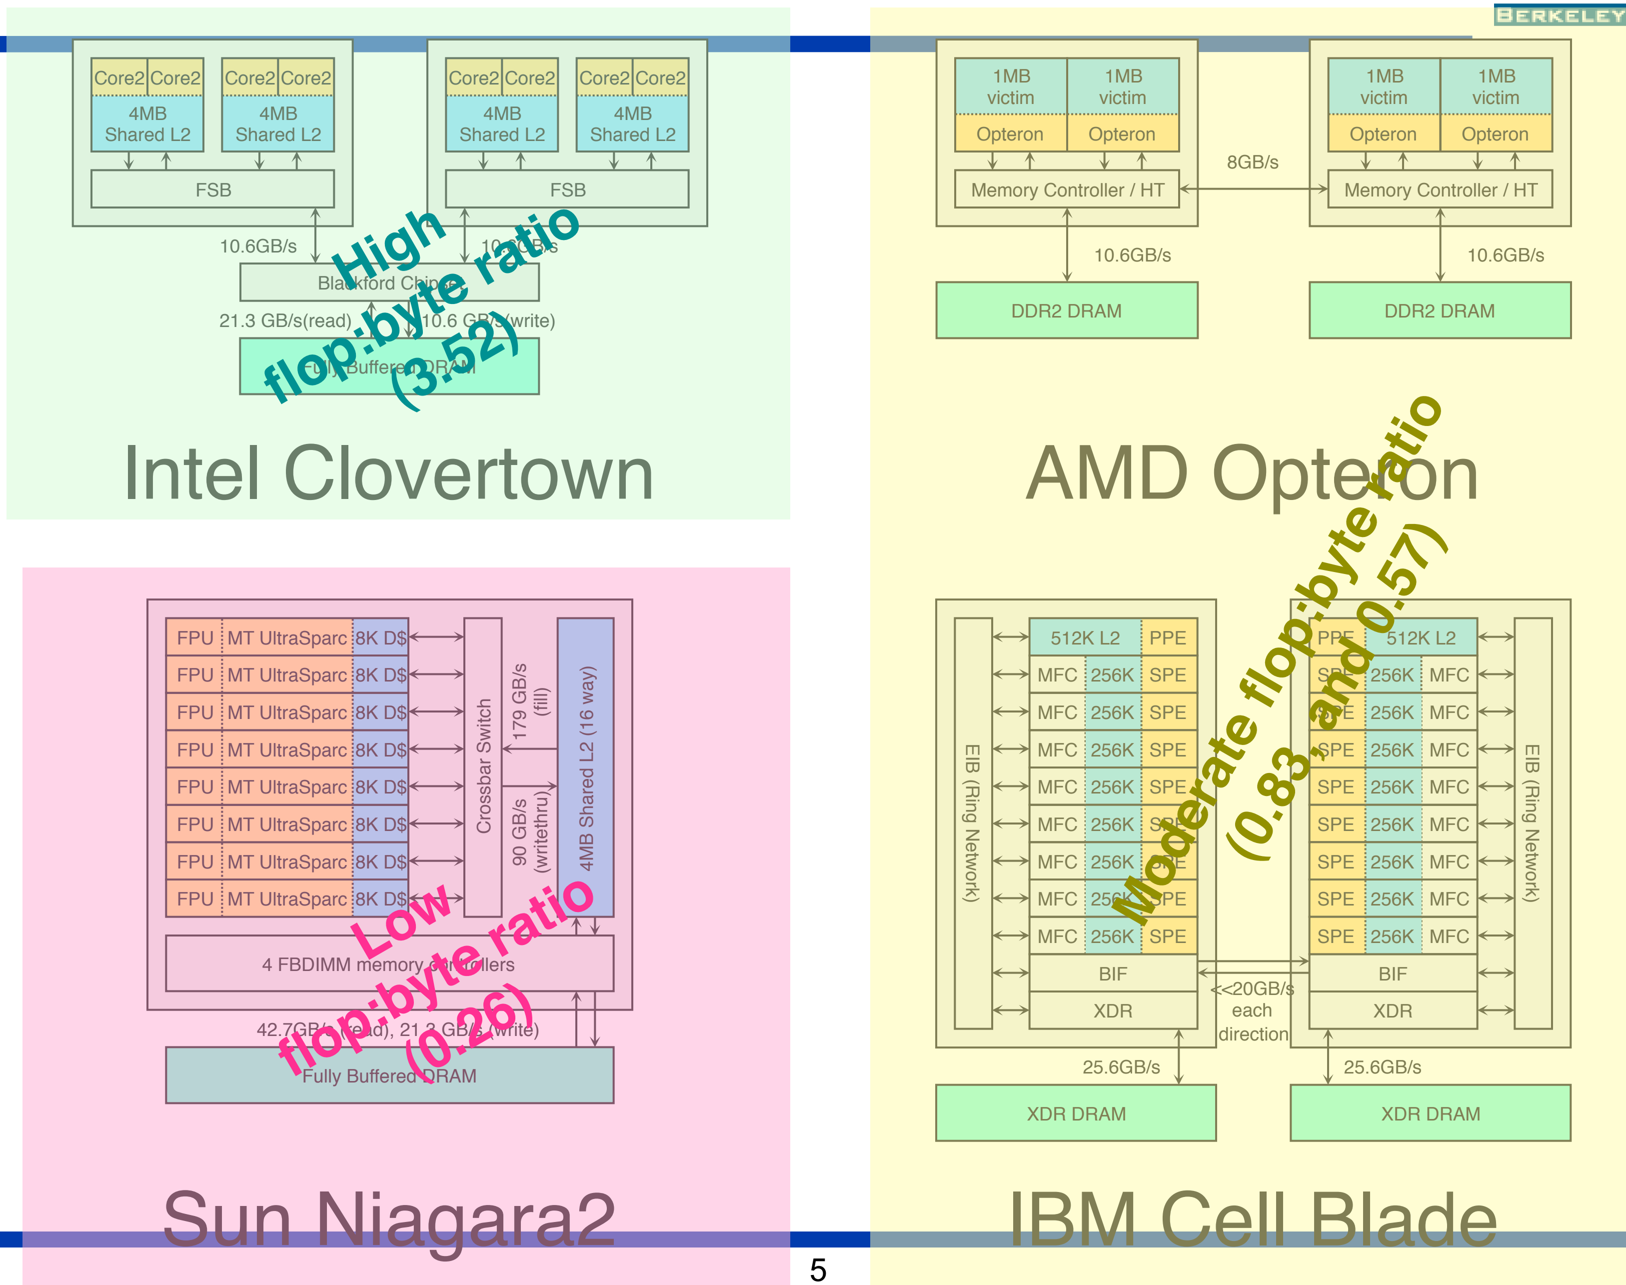

Four Multi-Core SMP Systems

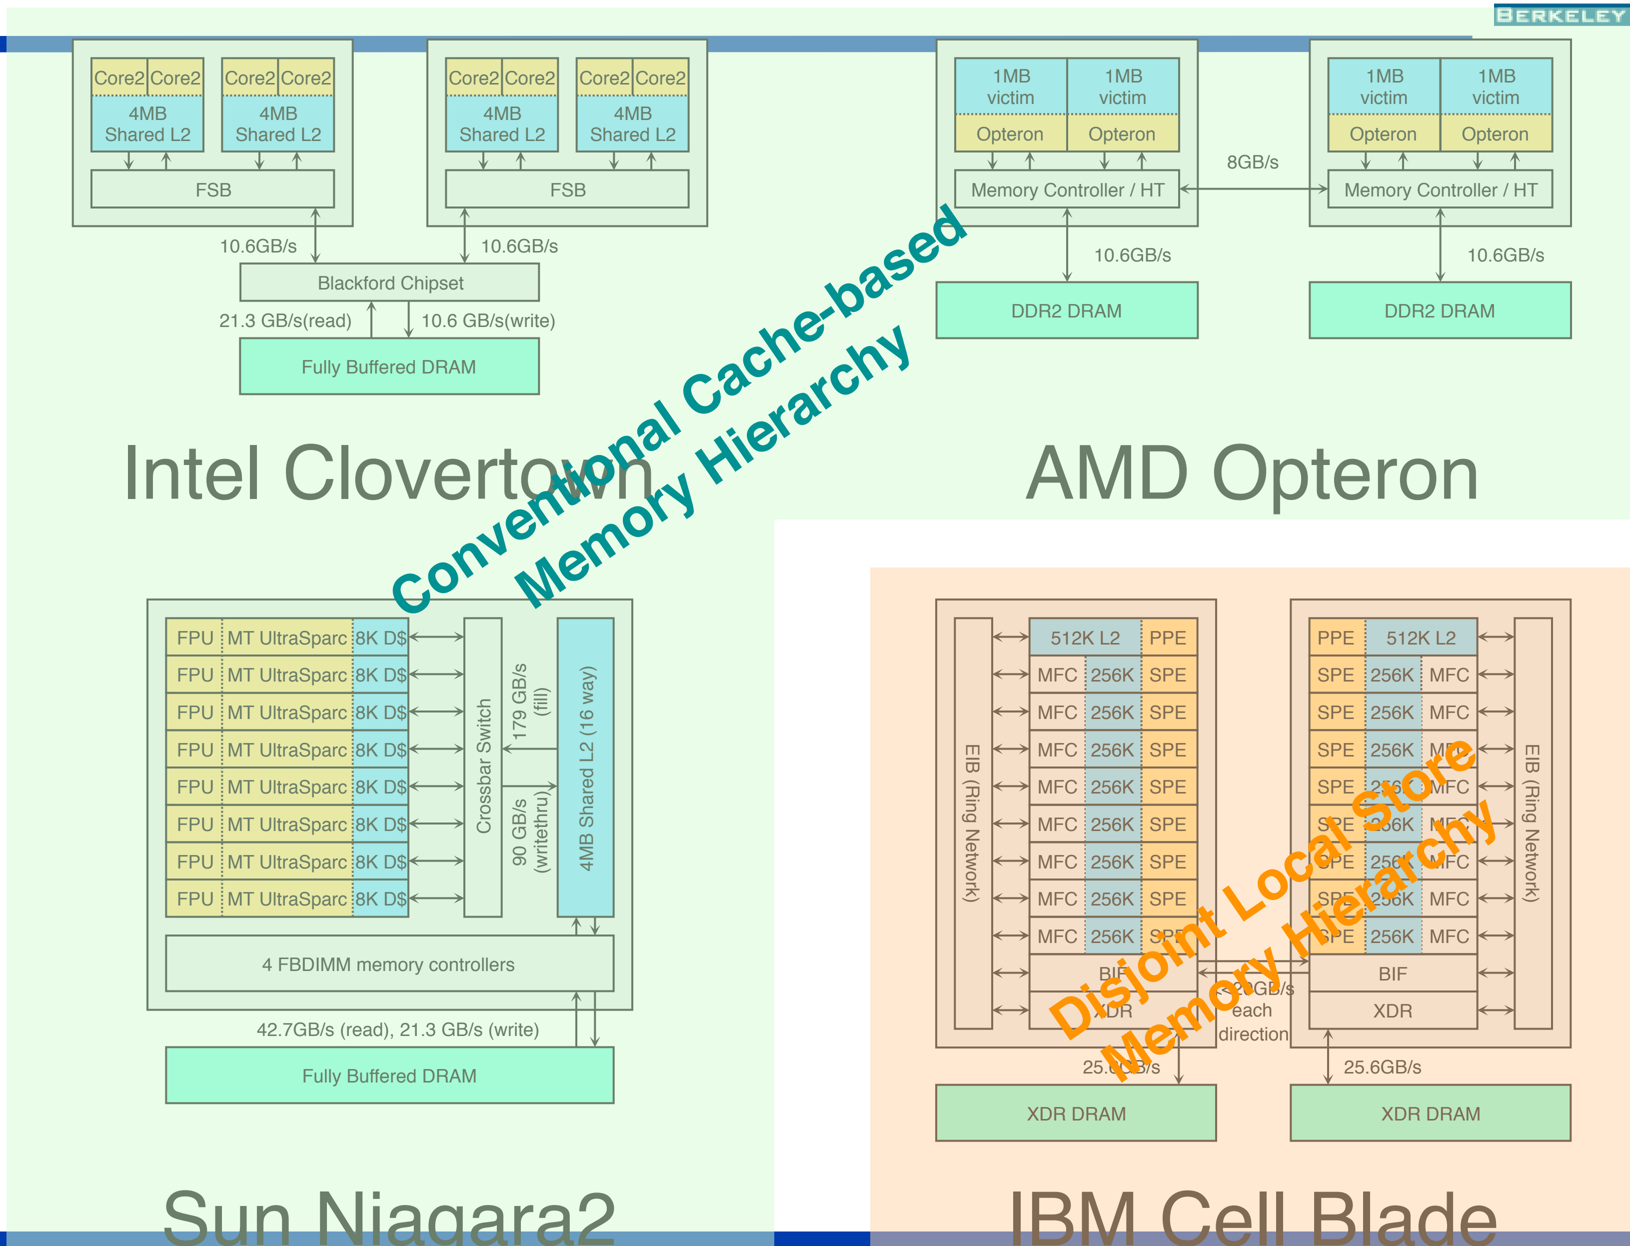

Four Multi-Core SMP Systems

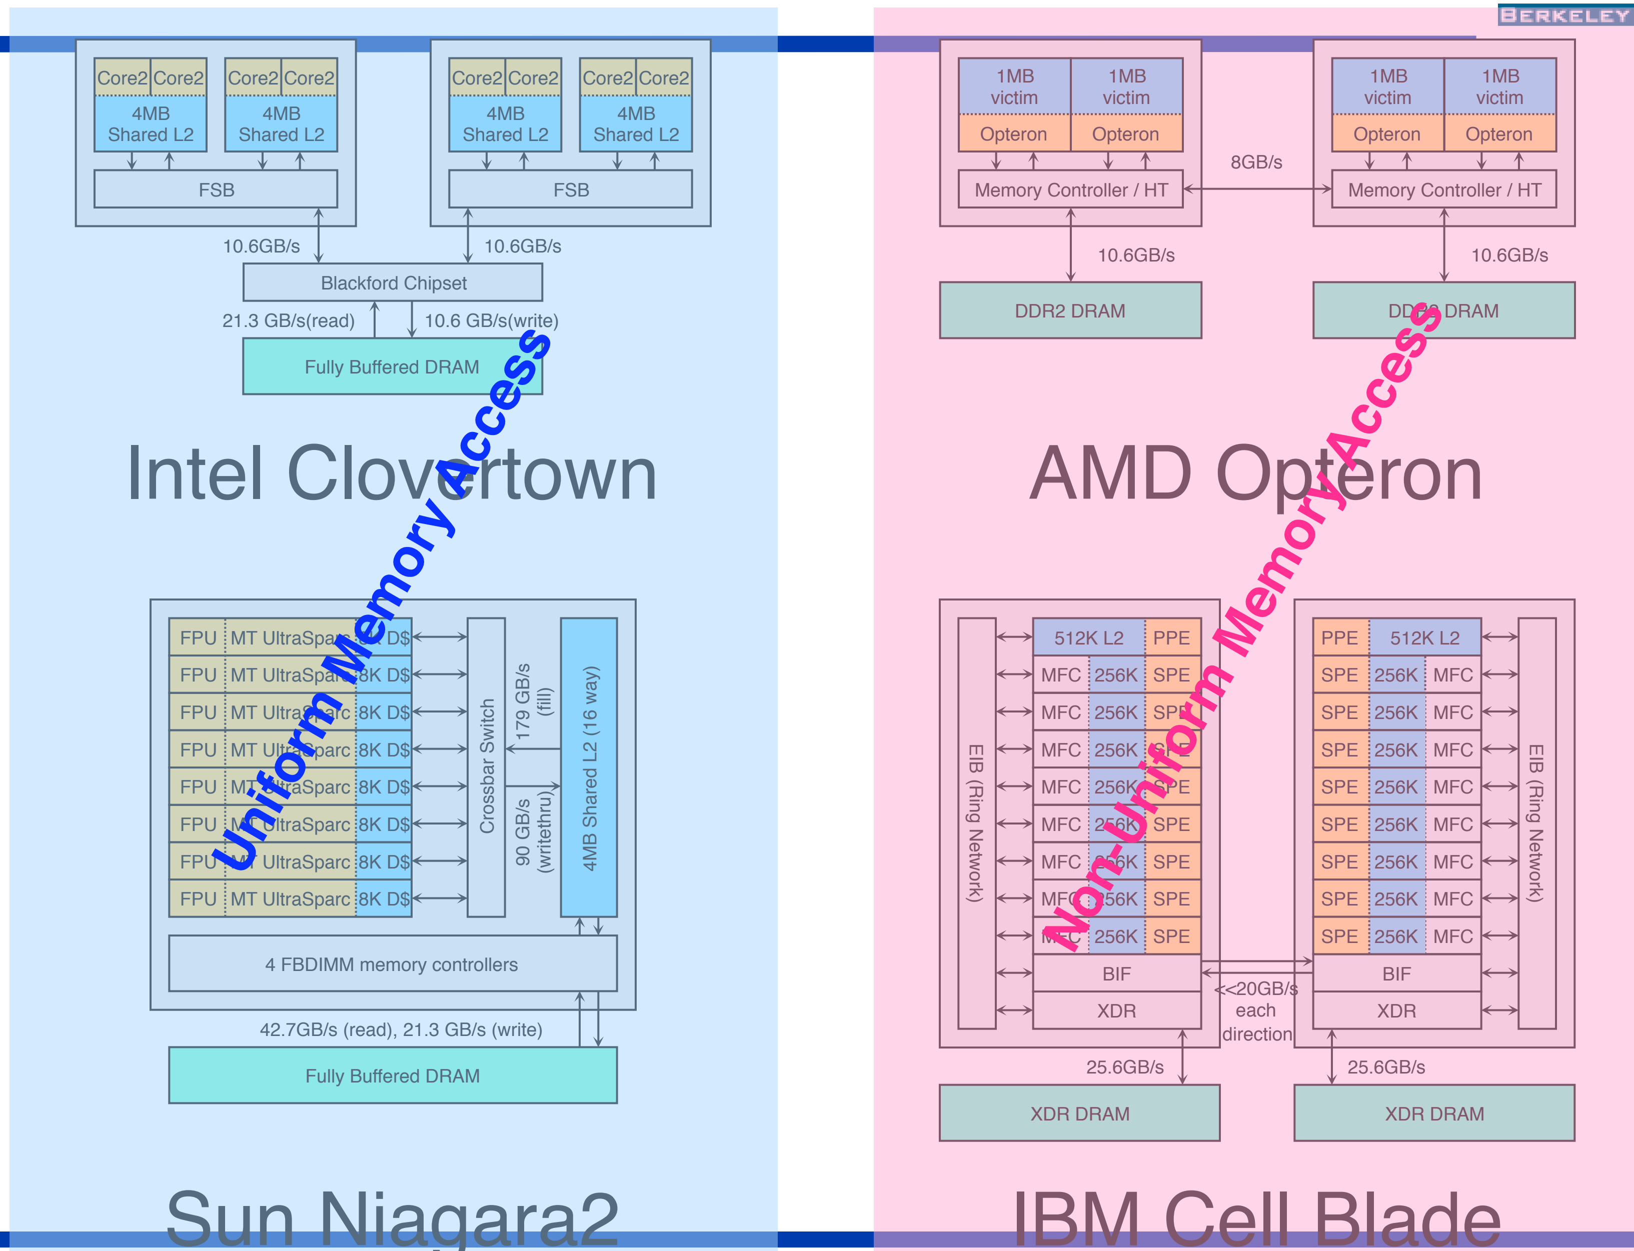

Four Multi-Core SMP Systems

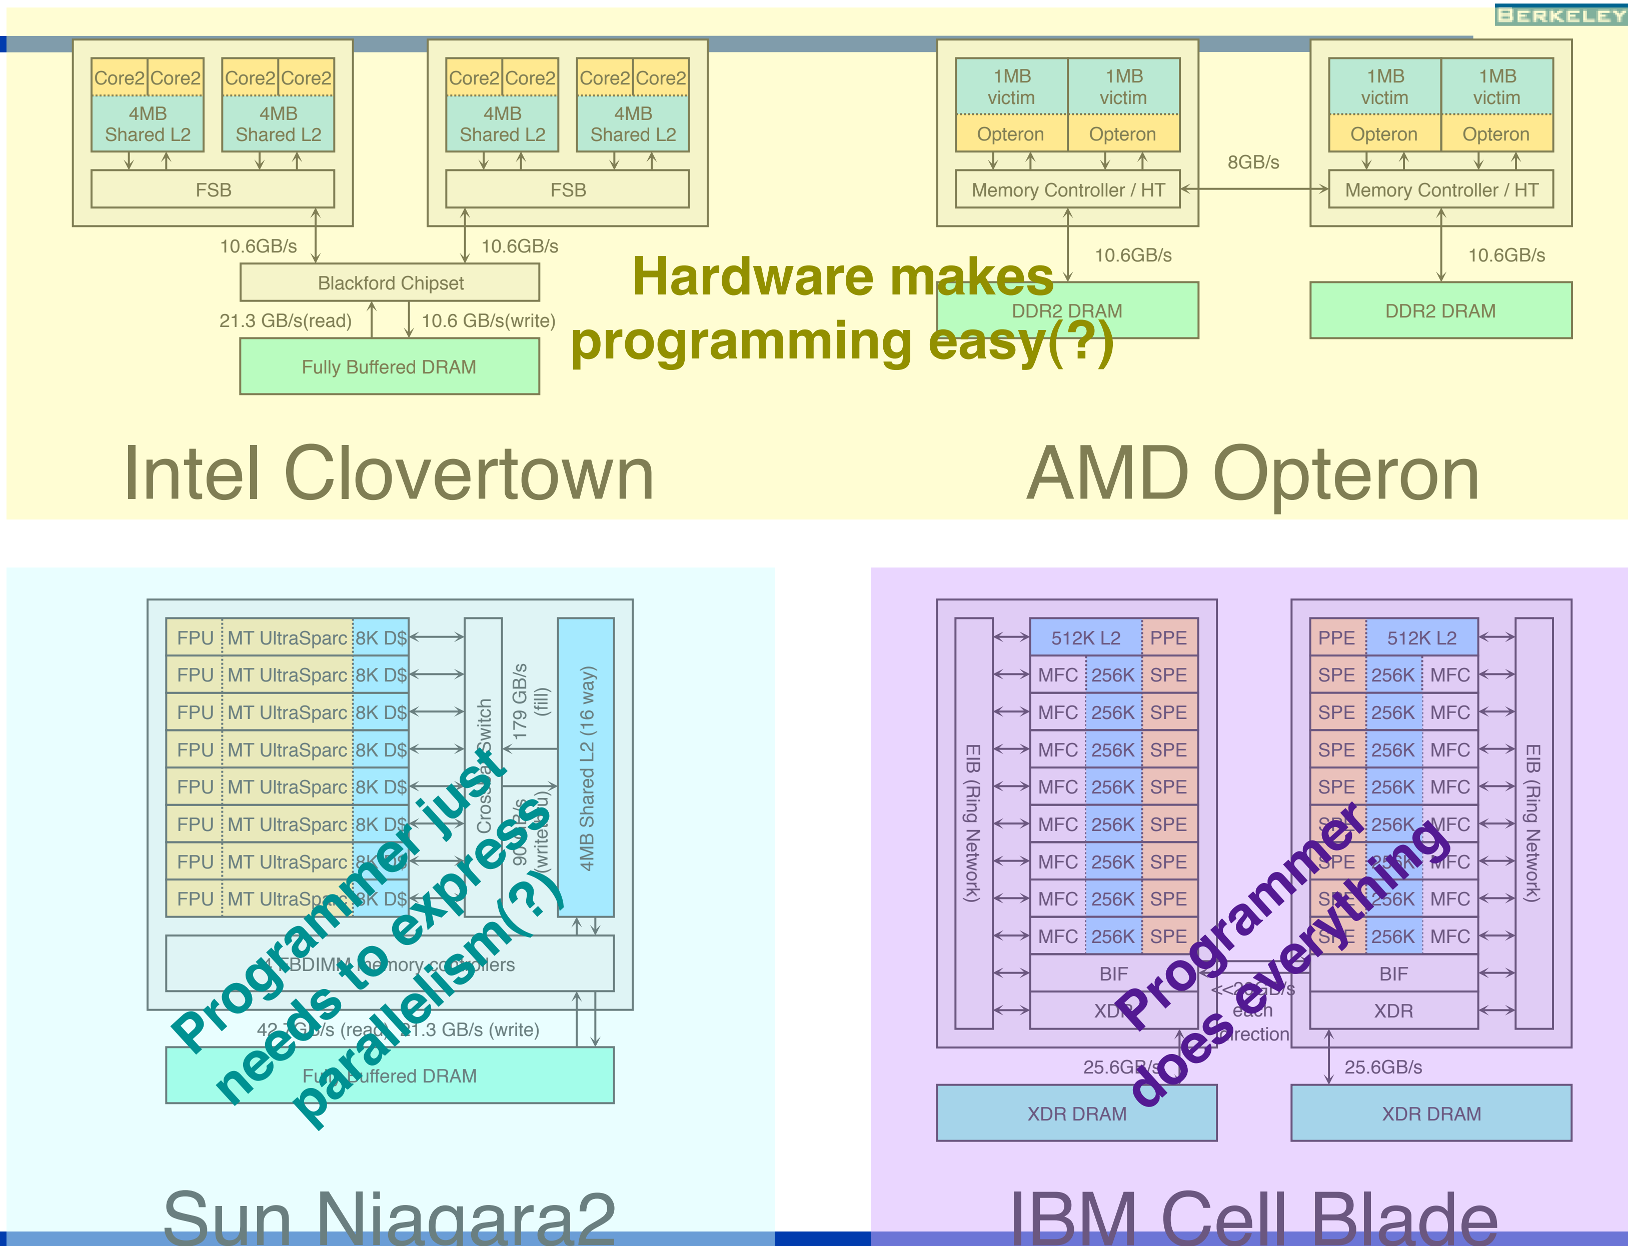

Naive Single Thread Performance

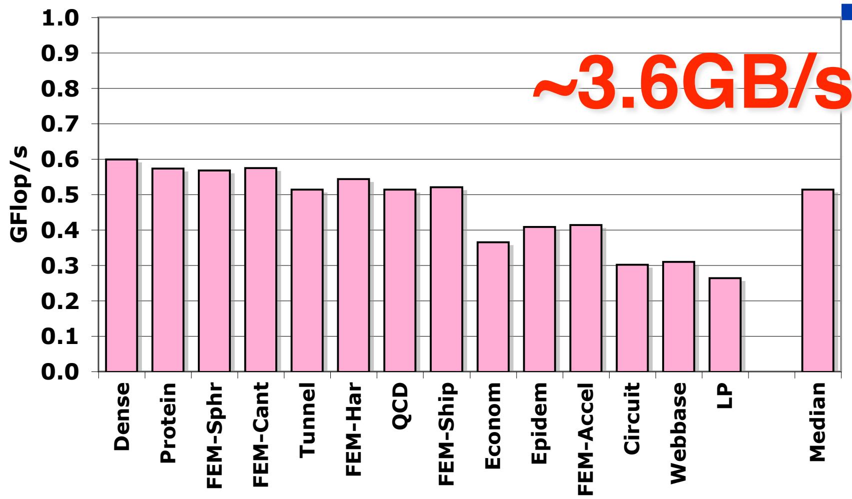

Intel Clovertown

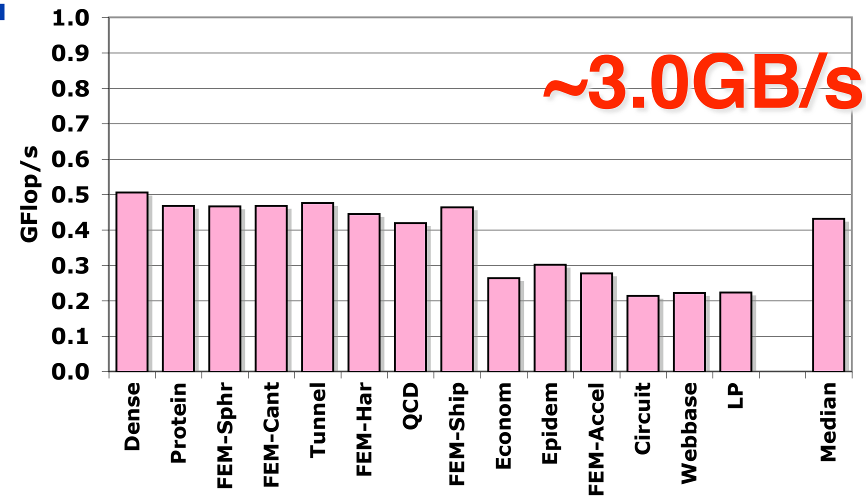

AMD Opteron

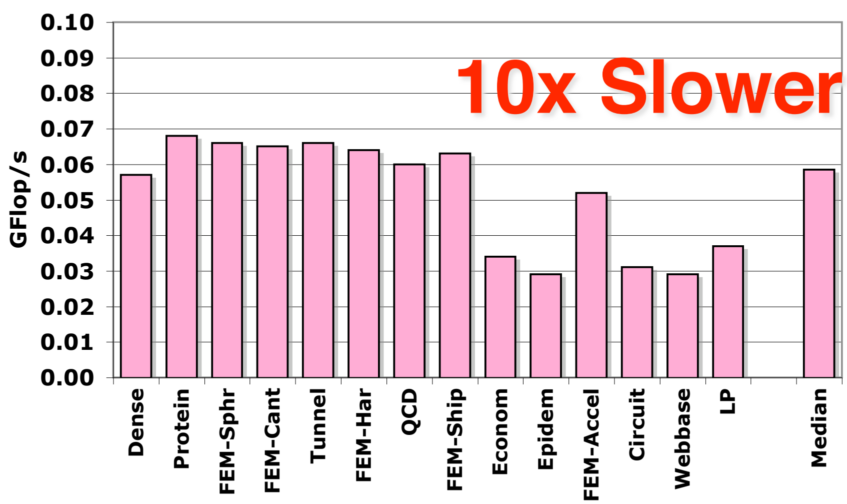

Sun Niagara2

Naive Single Thread

Naive Parallel Performance

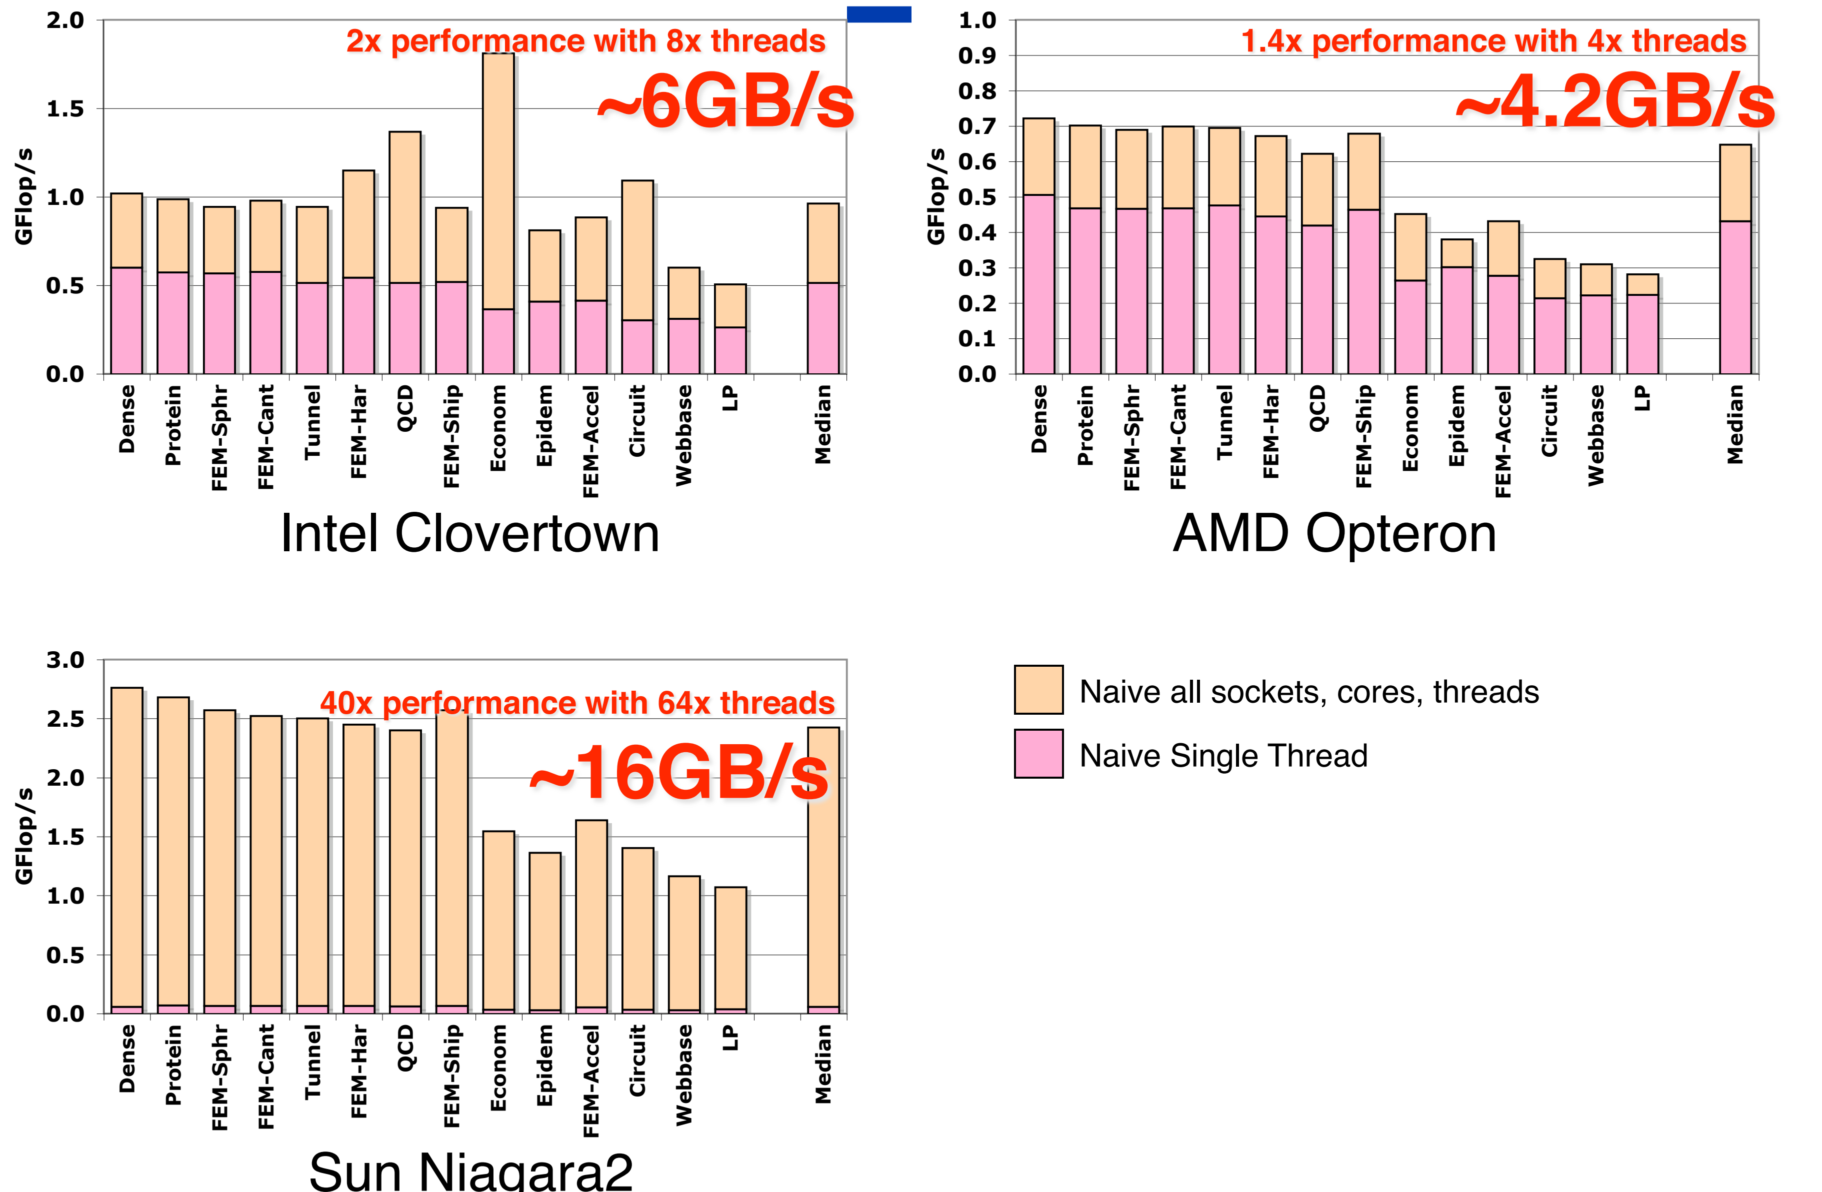

Exploiting NUMA / Affinity

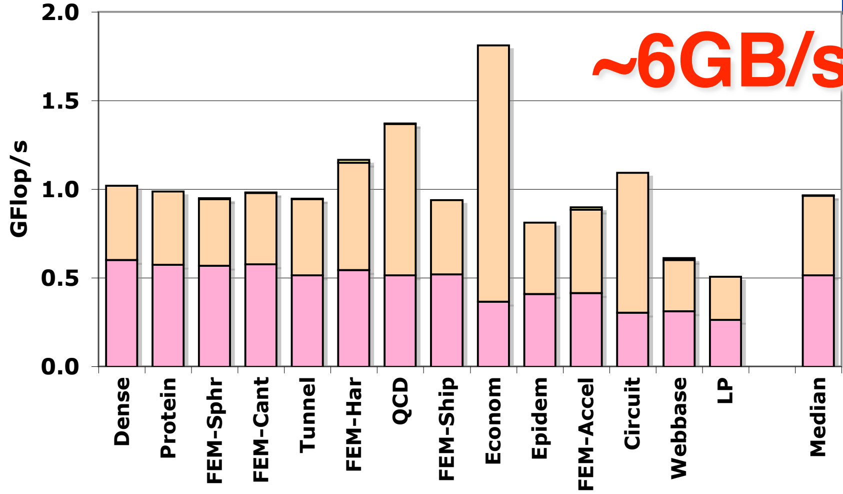

Intel Clovertown

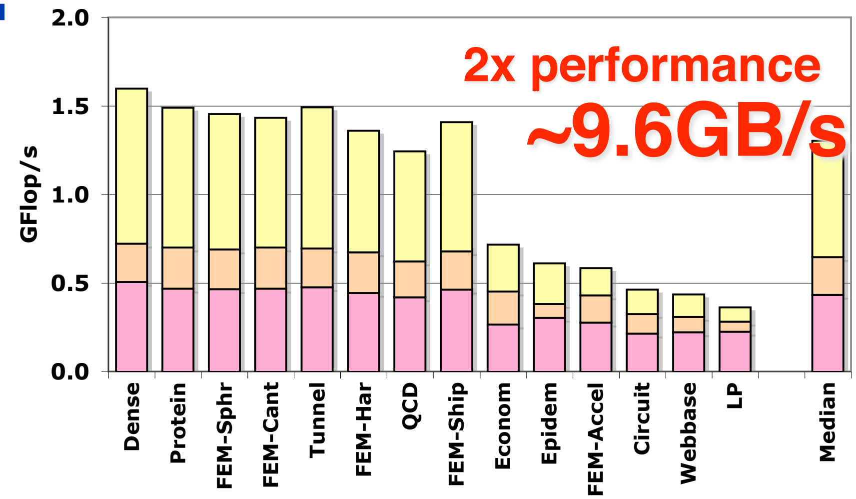

AMD Opteron

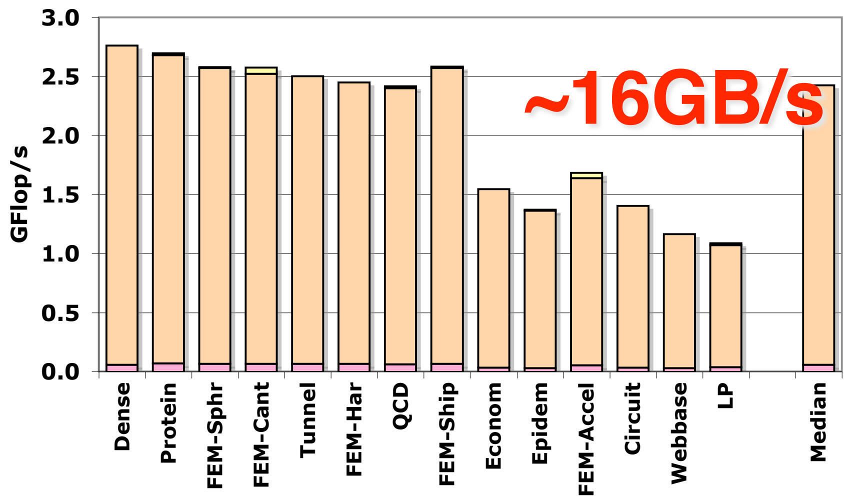

Sun Niagara2

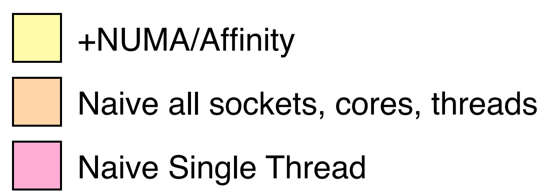

Exploiting Software Prefetch

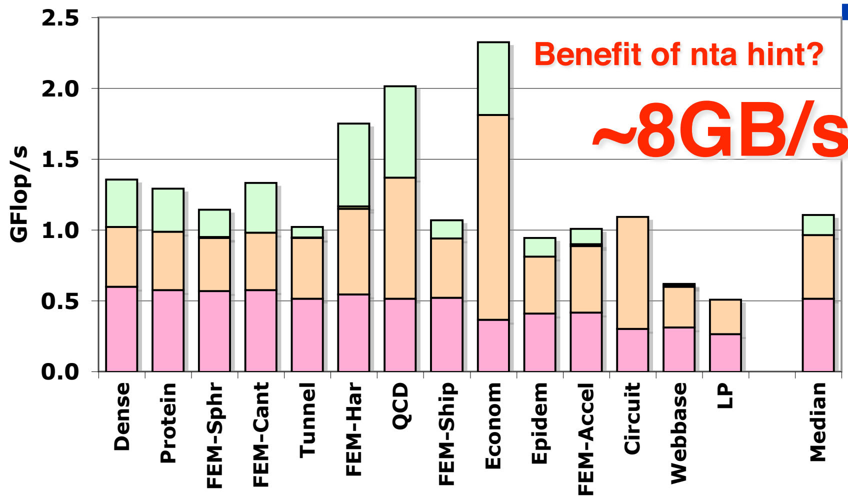

Intel Clovertown

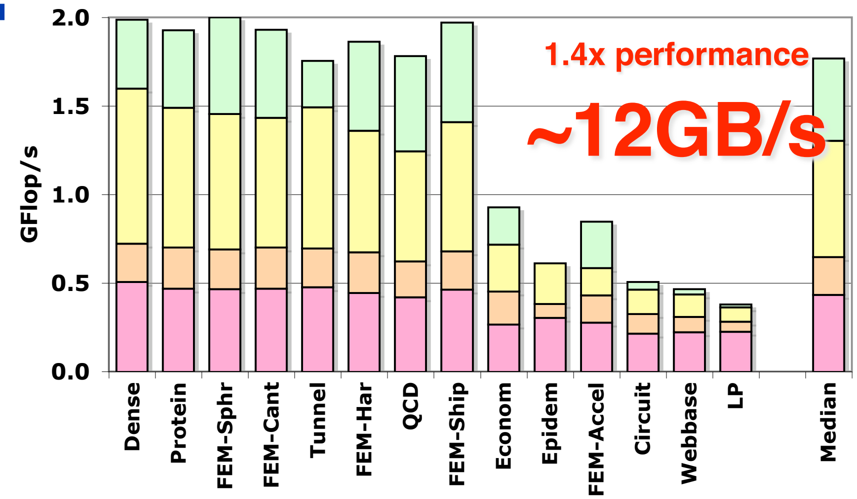

AMD Opteron

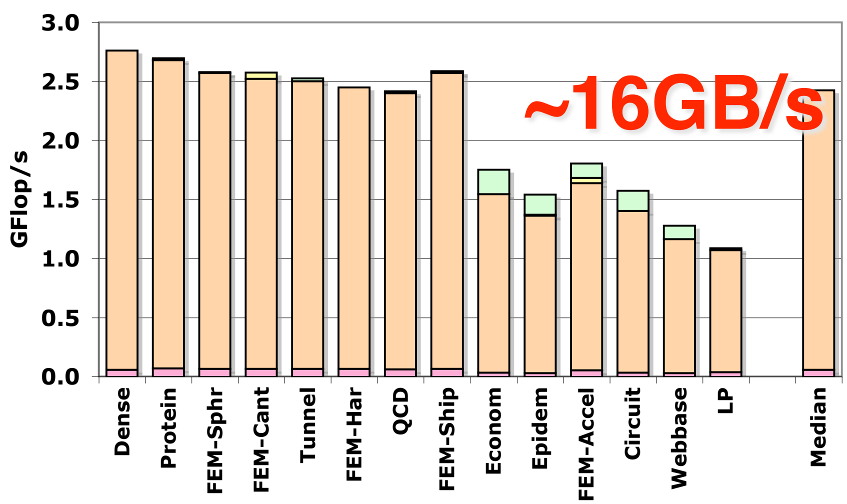

Sun Niagara2

- +Software Prefetching
- +NUMA/Affinity
- Naive all sockets, cores, threads
- Naive Single Thread

Exploiting Memory Traffic Minimization

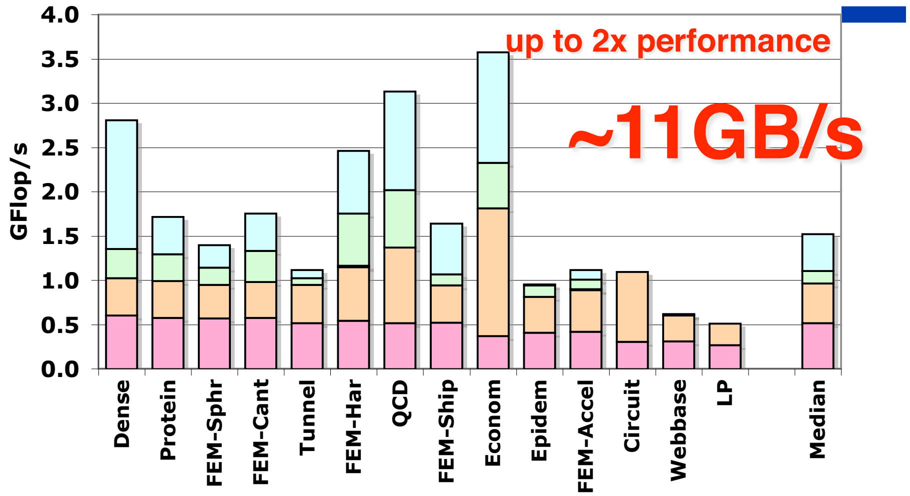

Intel Clovertown

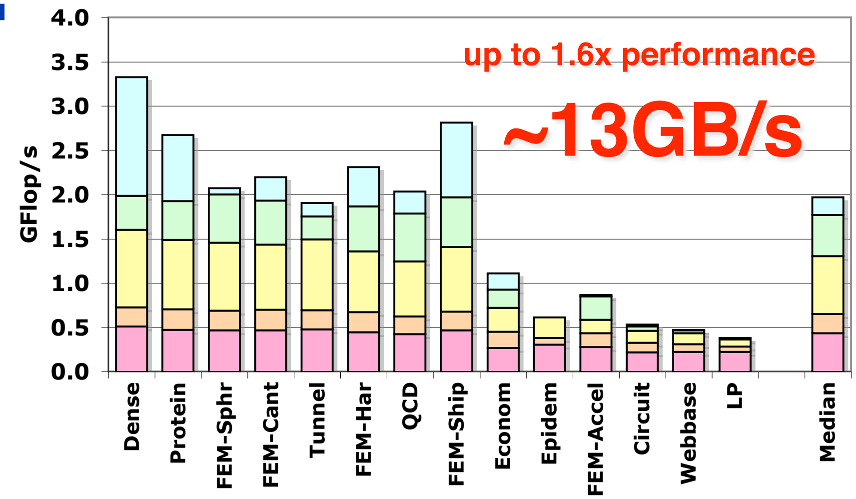

AMD Opteron

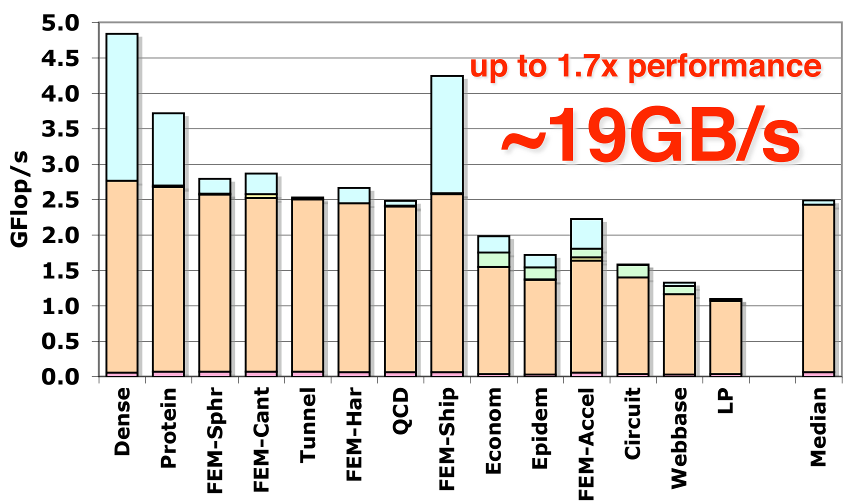

Sun Niagara2

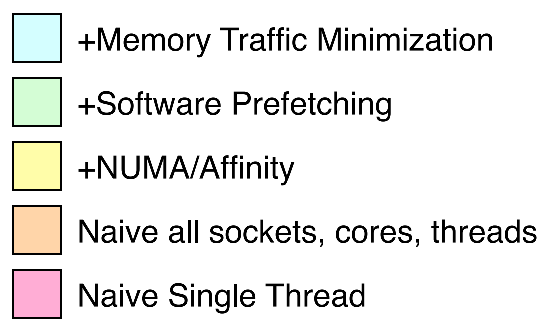

Berkeley Lab

Exploiting Cache Blocking

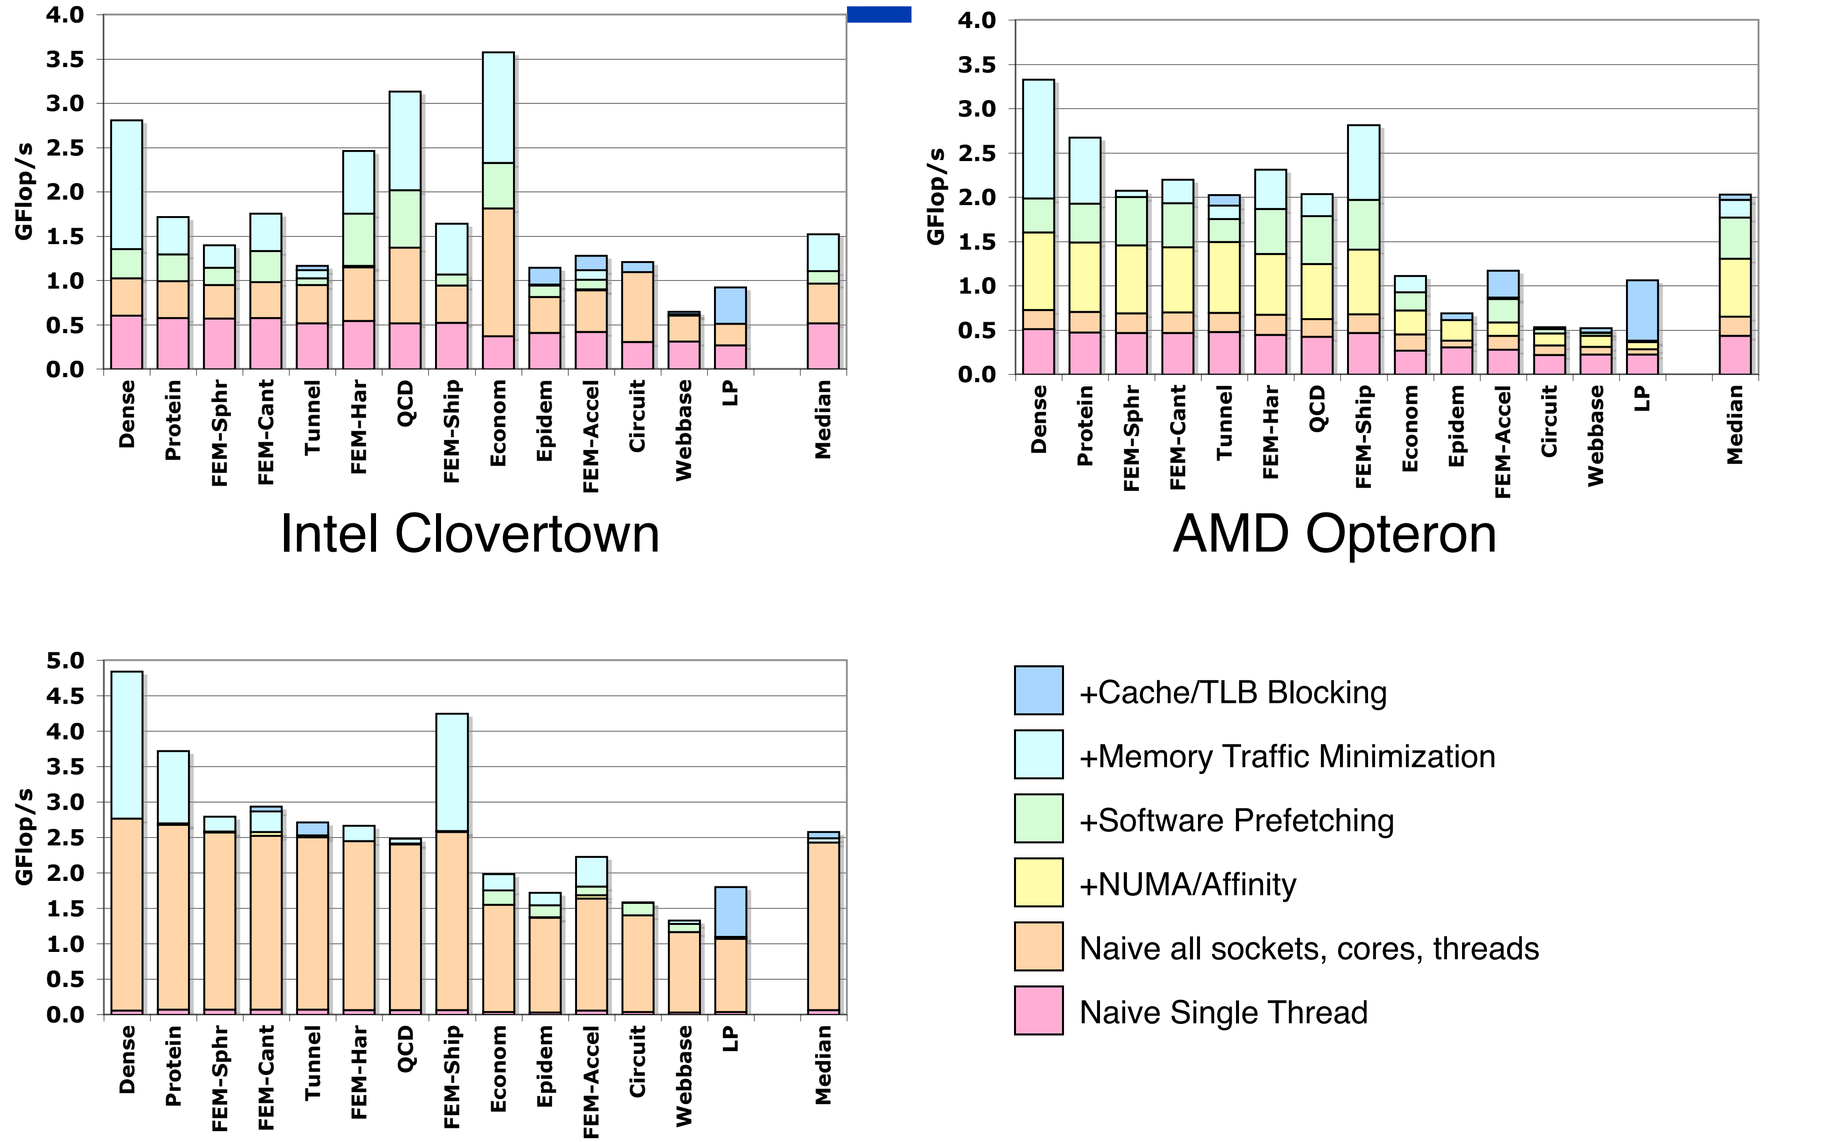

Intel Clovertown

AMD Opteron

Sun Niagara2

Exploiting More DIMMs/Ranks

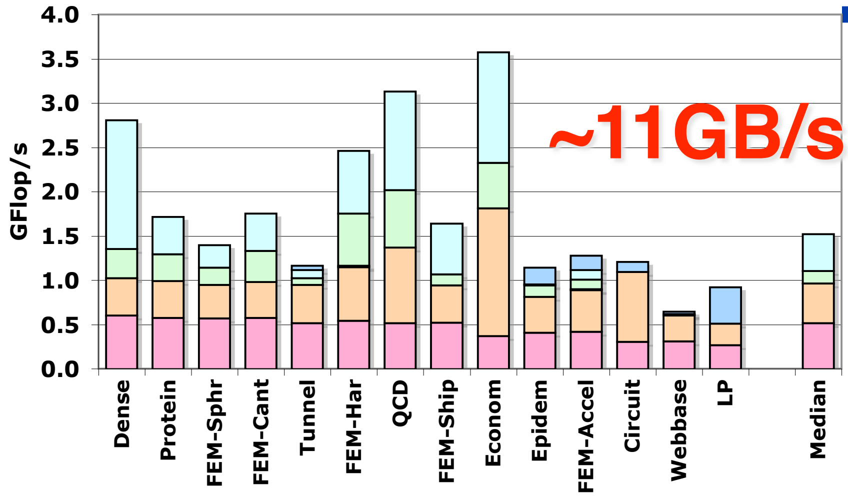

Intel Clovertown

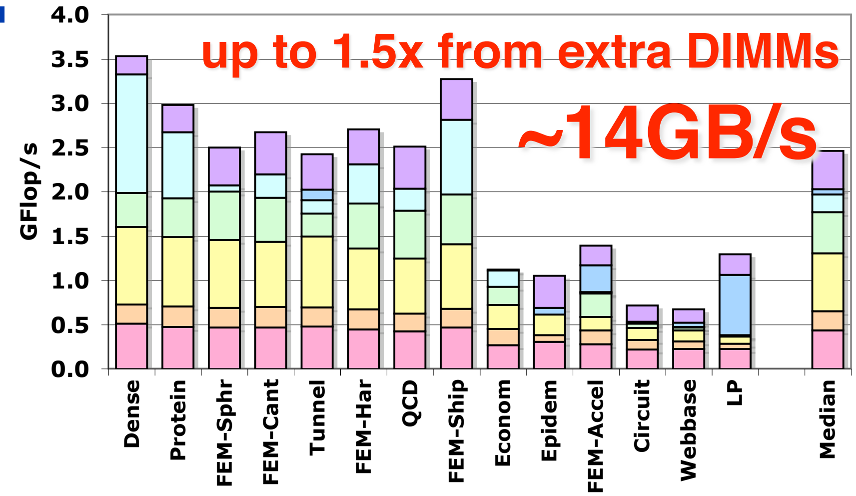

AMD Opteron

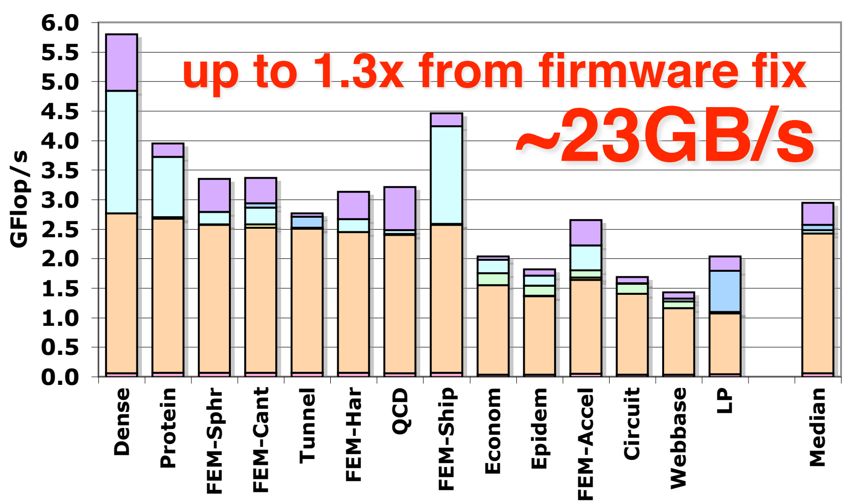

Sun Niagara2

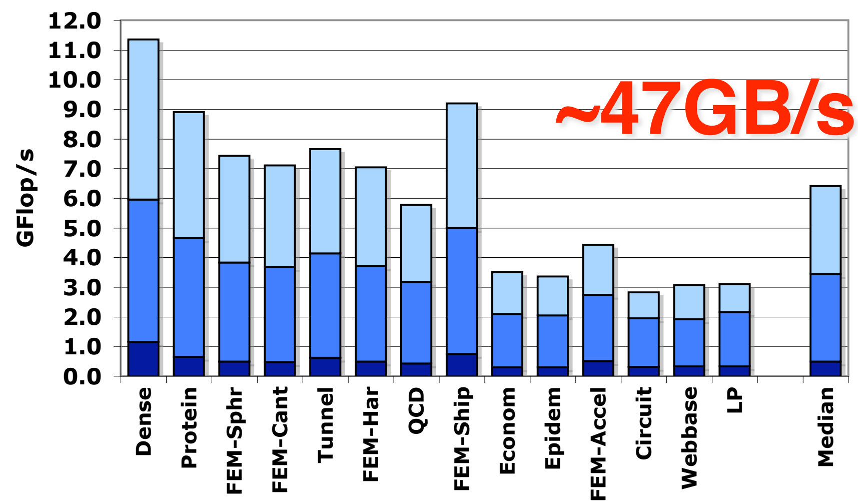

IBM Cell Blade

Conclusions of Auto-Tuning Study

- Naive implementations of sparse matrix-vector calculations on current multicore systems exhibit poor to mediocre performance, even with state-of-the-art compilers.
- Several types of code modifications, which can be done semi-automatically, substantially improve performance.
- Some changes require sophisticated search strategies and test runs.
- Optimal code alterations are different for each architecture.

Development of effective, broad-spectrum semi-automatic tuning techniques is a major research priority for the next 5-10 years.

- Semi-automatic tuning is a key component of the SciDAC Performance Engineering Research Institute (PERI).

In the meantime, performance for many applications on multicore processors will be decidedly suboptimal.

On the Bright Side: Some Applications Will Do Quite Well

- "Embarrassingly parallel" applications – e.g., running a large ensembles of individual instances of an application with different parameters.
- Applications that spawn independent tasks and combine results at the end, e.g., the "MapReduce" class in the Berkeley "dwarf" kernels.
- Computations that utilize highly tuned cache-aware libraries, such as LAPACK and ScaLAPACK, and FFTW.
- Applications employing high-precision arithmetic – double-double (32 digits), quad-double (64 digits), or arbitrary precision arithmetic.
  - These calculations enjoy very favorable data locality, and parallelization of the higher-level application is generally straightforward.

Applications of High-Precision Arithmetic in Modern Scientific Computing

- Highly nonlinear computations.
- Computations involving highly ill-conditioned linear systems.
- Computations involving data with very large dynamic range.
- Large computations on highly parallel computer systems.
- Computations where numerical sensitivity is not currently a major problem, but periodic testing is needed to ensure that results are reliable.
- Research problems in mathematics and mathematical physics that involve constant recognition and integer relation detection.

Few physicists, chemists or engineers are highly expert in numerical analysis. Thus high-precision arithmetic is often a better remedy for severe numerical round-off error, even if the error could, in principle, be improved with more advanced algorithms or coding techniques.

Available High-Precision Facilities

Vendor-supported arithmetic:

Type	Total Bits	Significant Digits	Support
IEEE Double	64	16	In hardware on almost all systems.
IEEE Extended	80	18	In hardware on Intel and AMD systems.
IEEE Quad	128	33	In software from some vendors (50-100X slower than IEEE double).

Non-commercial (free) software:

Type	Total Bits	Significant Digits	Support
Double-double	128	32	DDFUN90, QD.
Quad-double	256	64	QD.
Arbitrary	Any	Any	ARPREC, MPFUN90, GMP, MPFR.

Commercial software: *Mathematica, Maple*.

LBNL’s High-Precision Software

- QD: double-double (31 digits) and quad-double (62 digits).
- ARPREC: arbitrary precision.
- Low-level routines written in C++.
- C++ and Fortran-90 translation modules permit use with existing C++ and Fortran-90 programs -- only minor code changes are required.
- Includes many common functions: sqrt, cos, exp, gamma, etc.
- PSLQ, root finding, numerical integration.

Available at: **http://www.experimentalmath.info**

Authors: Xiaoye Li, Yozo Hida, Brandon Thompson and DHB

Some Real-World Applications of High-Precision Arithmetic

- Supernova simulations (32 or 64 digits).
- Climate modeling (32 digits).
- Planetary orbit calculations (32 digits).
- Coulomb *n*-body atomic system simulations (32-120 digits).
- Schrodinger solutions for lithium and helium atoms (32 digits).
- Electromagnetic scattering theory (32-100 digits).
- Studies of the fine structure constant of physics (32 digits).
- Scattering amplitudes of quarks, gluons and bosons (32 digits).
- Theory of nonlinear oscillators (64 digits).

DHB, "High-Precision Arithmetic in Scientific Computation," *Computing in Science and Engineering*, May-Jun, 2005, pg. 54-61.
DHB and Jonathan M. Borwein, "High-Precision Computation and Mathematical Physics," *XII Advanced Computing and Analysis Techniques in Physics Research*, 2008, to appear, http://crd.lbl.gov/~dhbailey/dhbpapers/numerical-bugs.pdf.

Supernova Simulations

- Researchers at LBNL are using QD to solve for non-local thermodynamic equilibrium populations of iron and other atoms in the atmospheres of supernovas.
- Iron may exist in several species, so it is necessary to solve for all species simultaneously.
- Since the relative population of any state from the dominant state is proportional to the exponential of the ionization energy, the dynamic range of these values can be very large.
- The quad-double portion now dominates the entire computation.

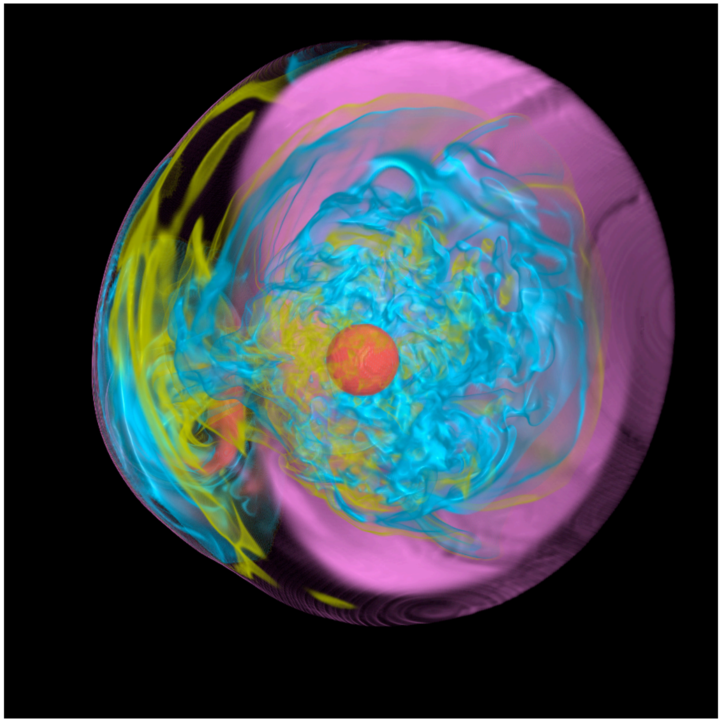

P. H. Hauschildt and E. Baron, "The Numerical Solution of the Expanding Stellar Atmosphere Problem," *Journal Computational and Applied Mathematics*, vol. 109 (1999), pg. 41-63.

Climate Modeling

- Climate and weather simulations are fundamentally chaotic – if microscopic changes are made to the current state, soon the future state is quite different.
- In practice, computational results are altered even if minor changes are made to the code or the system.
- This numerical variation is a major nuisance for code maintenance.
- He and Ding of LBNL found that by using double-double arithmetic to implement a key inner product loop, most of this numerical variation disappeared.

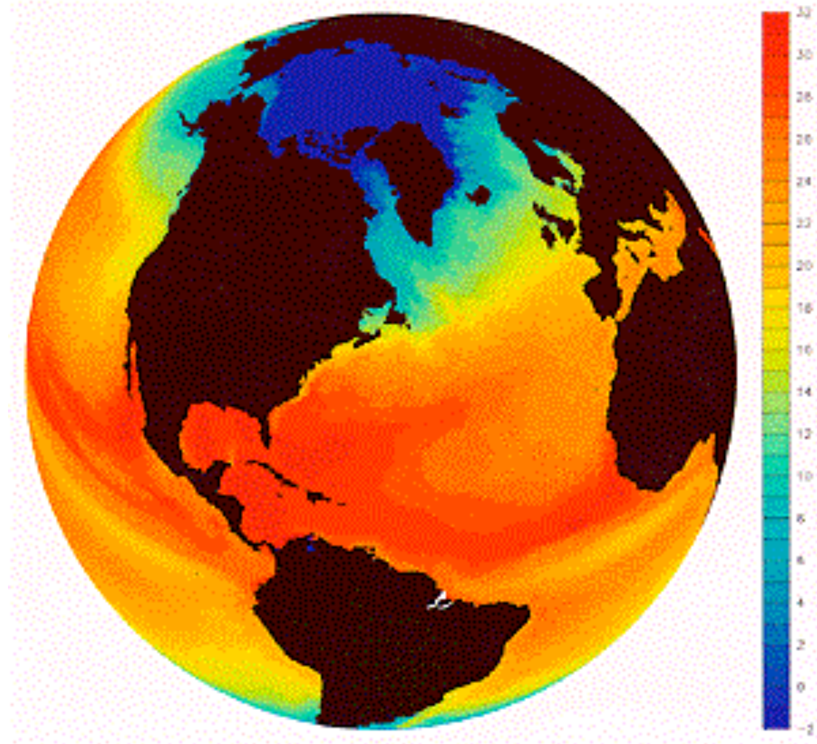

Yun He and Chris Ding, "Using Accurate Arithmetics to Improve Numerical Reproducibility and Stability in Parallel Applications," *Journal of Supercomputing*, vol. 18, no. 3 (Mar 2001), pg. 259-277.

Coulomb N-Body Atomic System Simulations

- Alexei Frolov of Queen's University in Canada has used MPFUN90 to solve a generalized eigenvalue problem that arises in Coulomb n-body interactions.
- Matrices are typically 5,000 x 5,000 and are very nearly singular.
- Frolov has also computed elements of the Hamiltonian matrix and the overlap matrix in four- and five-body systems.
- These computations typically require 120-digit arithmetic.

"We can consider and solve the bound state few-body problems which have been beyond our imagination even four years ago." – Frolov

A. M. Frolov and DHB, "Highly Accurate Evaluation of the Few-Body Auxiliary Functions and Four-Body Integrals," *Journal of Physics B*, vol. 36, no. 9 (14 May 2003), pg. 1857-1867.

The PSLQ Integer Relation Algorithm

Let (x_n) be a given vector of real numbers. An integer relation algorithm finds integers (a_n) such that

$$a_1 x_1 + a_2 x_2 + \cdots + a_n x_n = 0$$

(or within "epsilon" of zero, where epsilon = 10^{-p} and p is the precision).

At the present time the "PSLQ" algorithm of mathematician-sculptor Helaman Ferguson is the most widely used integer relation algorithm. It was named one of ten "algorithms of the century" by *Computing in Science and Engineering*.

PSLQ (or any other integer relation scheme) requires very high precision (at least $n*d$ digits, where d is the size in digits of the largest a_k), both in the input data and in the operation of the algorithm.

1. H. R. P. Ferguson, DHB and S. Arno, "Analysis of PSLQ, An Integer Relation Finding Algorithm," *Mathematics of Computation*, vol. 68, no. 225 (Jan 1999), pg. 351-369.
2. DHB and D. J. Broadhurst, "Parallel Integer Relation Detection: Techniques and Applications," *Mathematics of Computation*, vol. 70, no. 236 (Oct 2000), pg. 1719-1736.

Decrease of $\log_{10}(\min |x_i|)$ in PSLQ

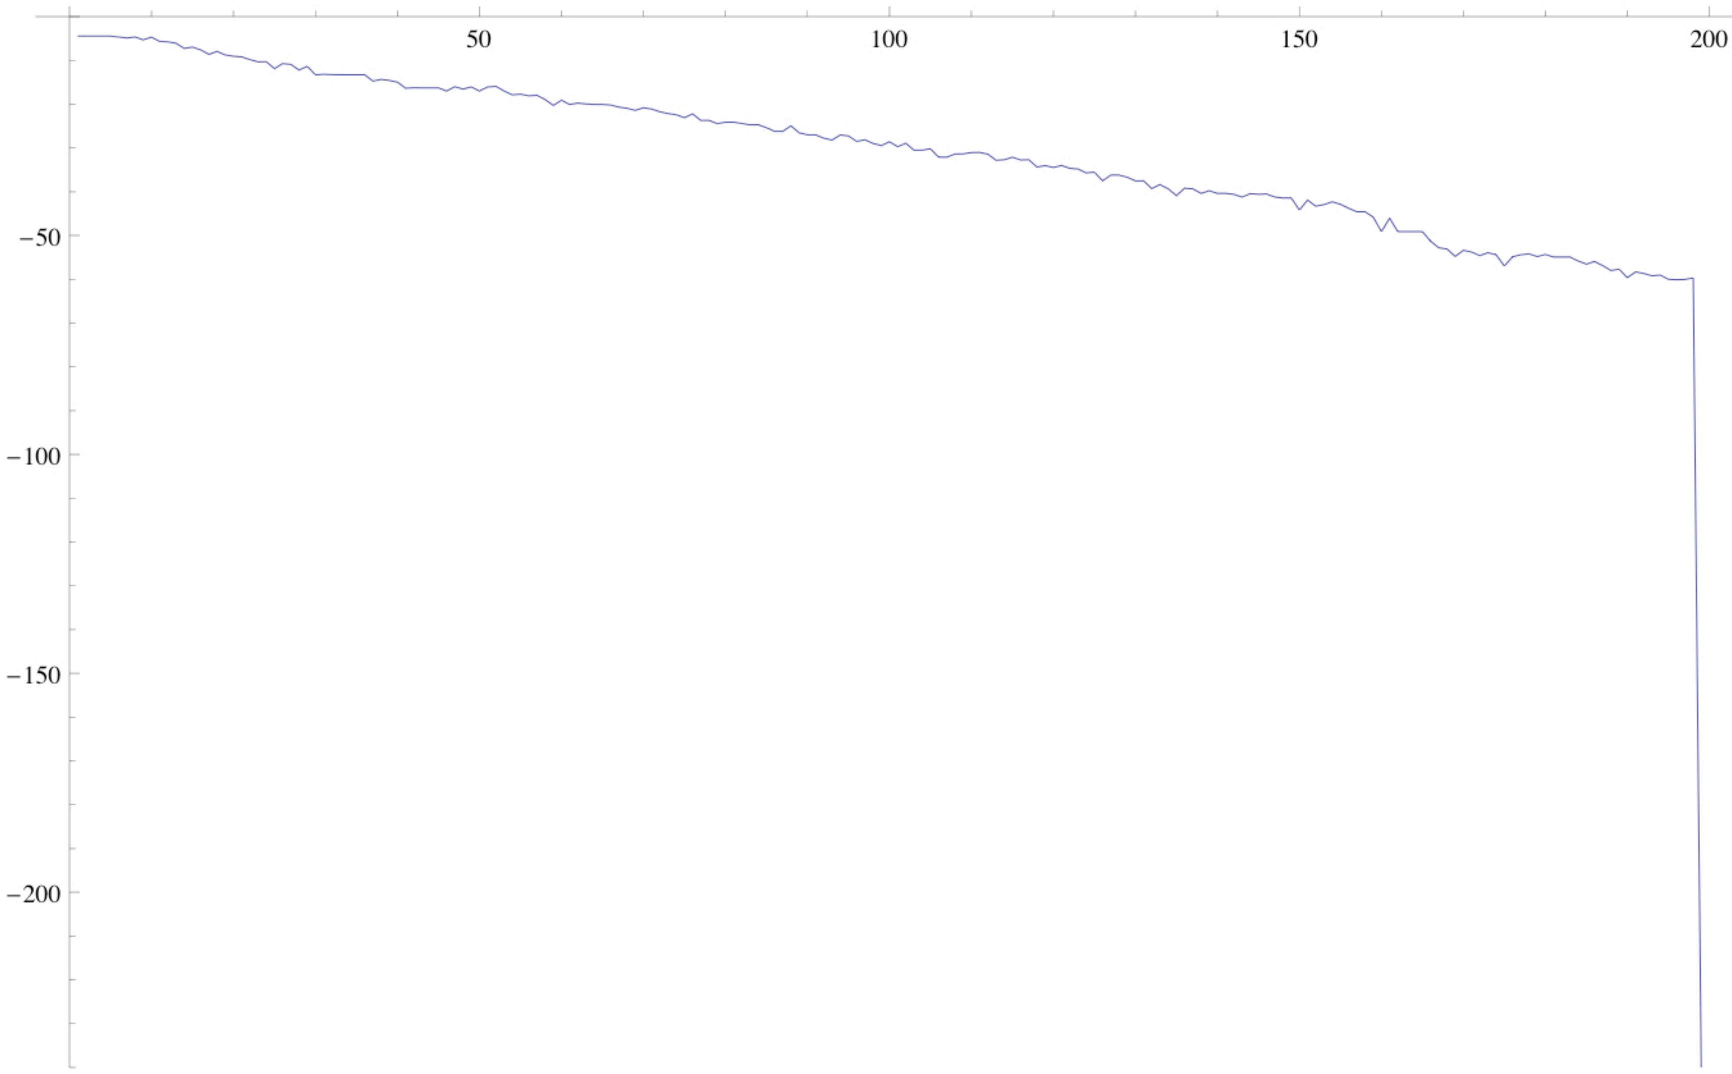

Bifurcation Points in Chaos Theory

Let $t = B_3$ = the smallest r such that the "logistic iteration"

$$x_{n+1} = r x_n (1 - x_n)$$

exhibits 8-way periodicity instead of 4-way periodicity.

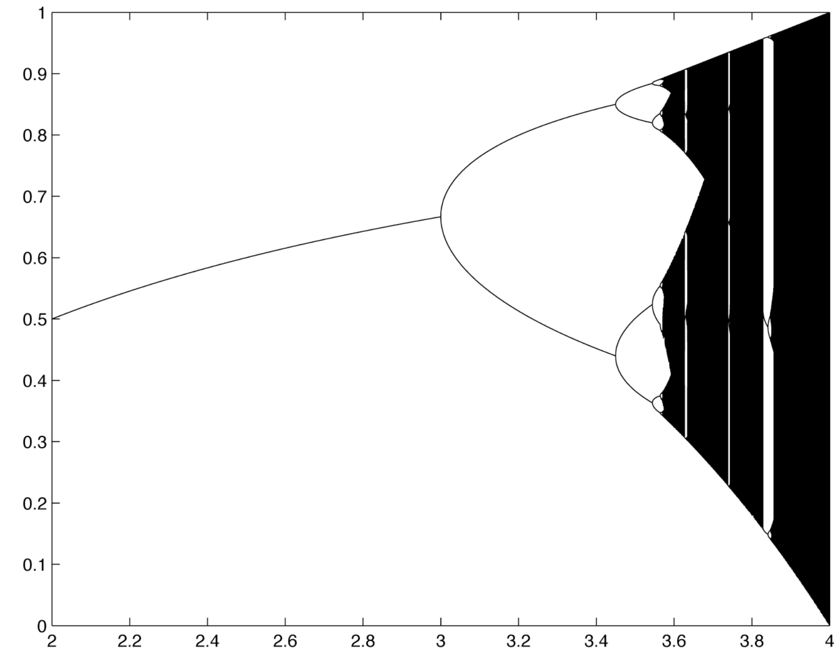

By means of a sequential approximation scheme, one can obtain the numerical value of t to any desired precision:

3.544090359551922853615965986604804540583099845444573675457813...

Applying PSLQ to $(1, t, t^2, t^3, \ldots, t^{12})$, we obtained the result that t is a root of the polynomial:

$$\begin{aligned} 0 &= 4913 + 2108t^2 - 604t^3 - 977t^4 + 8t^5 + 44t^6 + 392t^7 \\ &\quad -193t^8 - 40t^9 + 48t^{10} - 12t^{11} + t^{12} \end{aligned}$$

The BBP Formula for Pi

In 1996, at the suggestion of Peter Borwein, Simon Plouffe used DHB's PSLQ program and arbitrary precision software to discover this new formula for π:

$$\pi = \sum_{k=0}^{\infty} \frac{1}{16^k} \left(\frac{4}{8k+1} - \frac{2}{8k+4} - \frac{1}{8k+5} - \frac{1}{8k+6} \right)$$

This formula permits one to compute binary (or hexadecimal) digits of pi beginning at an arbitrary starting position, using a very simple scheme that can run on any system with standard 64-bit or 128-bit arithmetic.

Numerous other formulas of this type have since been found for various other mathematical constants.

1. DHB, P. B. Borwein and S. Plouffe, "On the Rapid Computation of Various Polylogarithmic Constants," *Mathematics of Computation*, vol. 66, no. 218 (Apr 1997), pg. 903-913.
2. J. M. Borwein, W. F. Galway and D. Borwein, "Finding and Excluding b-ary Machin-Type BBP Formulae," *Canadian Journal of Mathematics*, vol. 56 (2004), pg 1339-1342.
3. DHB, "A Compendium of BBP-Type Formulas," 2004, available at http://crd.lbl.gov/~dhbailey/dhbpapers/bbp-formulas.pdf.

Tanh-Sinh Quadrature

Given $f(x)$ defined on (-1,1), define $g(t) = \tanh(\pi/2 \sinh t)$. Then setting $x = g(t)$ yields

$$\int_{-1}^{1} f(x)\,dx \quad = \quad \int_{-\infty}^{\infty} f(g(t))g'(t)\,dt \;\approx\; h \sum_{j=-N}^{N} w_j f(x_j),$$

where $x_j = g(hj)$ and $w_j = g'(hj)$. Since $g'(t)$ goes to zero very rapidly for large t, the product $f(g(t))\, g'(t)$ typically is a nice bell-shaped function for which the Euler-Maclaurin formula implies that the simple summation above is remarkably accurate. Reducing h by half typically doubles the number of correct digits.

For our applications, we have found that tanh-sinh is the best general-purpose integration scheme for functions with vertical derivatives or singularities at endpoints, or for any function at very high precision (> 1000 digits). Otherwise we use Gaussian quadrature.

1. DHB, Xiaoye S. Li and Karthik Jeyabalan, "A Comparison of Three High-Precision Quadrature Schemes," *Experimental Mathematics*, vol. 14 (2005), no. 3, pg. 317-329.
2. H. Takahasi and M. Mori, "Double Exponential Formulas for Numerical Integration," Publications of RIMS, Kyoto University, vol. 9 (1974), pg. 721–741.

A Log-Tan Integral Identity

$$\frac{24}{7\sqrt{7}}\int_{\pi/3}^{\pi/2}\log\left|\frac{\tan t+\sqrt{7}}{\tan t-\sqrt{7}}\right|dt \quad = \quad L_{-7}(2) \quad =$$

$$\sum_{n=0}^{\infty}\left[\frac{1}{(7n+1)^2}+\frac{1}{(7n+2)^2}-\frac{1}{(7n+3)^2}+\frac{1}{(7n+4)^2}-\frac{1}{(7n+5)^2}-\frac{1}{(7n+6)^2}\right]$$

This identity arises from analysis of volumes of knot complements in hyperbolic space. This is simplest of 998 related identities.

We have verified this numerically to 20,000 digits, using a highly parallel tanh-sinh quadrature code, running on 1024 cores.

DHB, J. M. Borwein, V. Kapoor and E. Weisstein, "Ten Problems in Experimental Mathematics," *American Mathematical Monthly*, vol. 113, no. 6 (Jun 2006), pg. 481-409 .

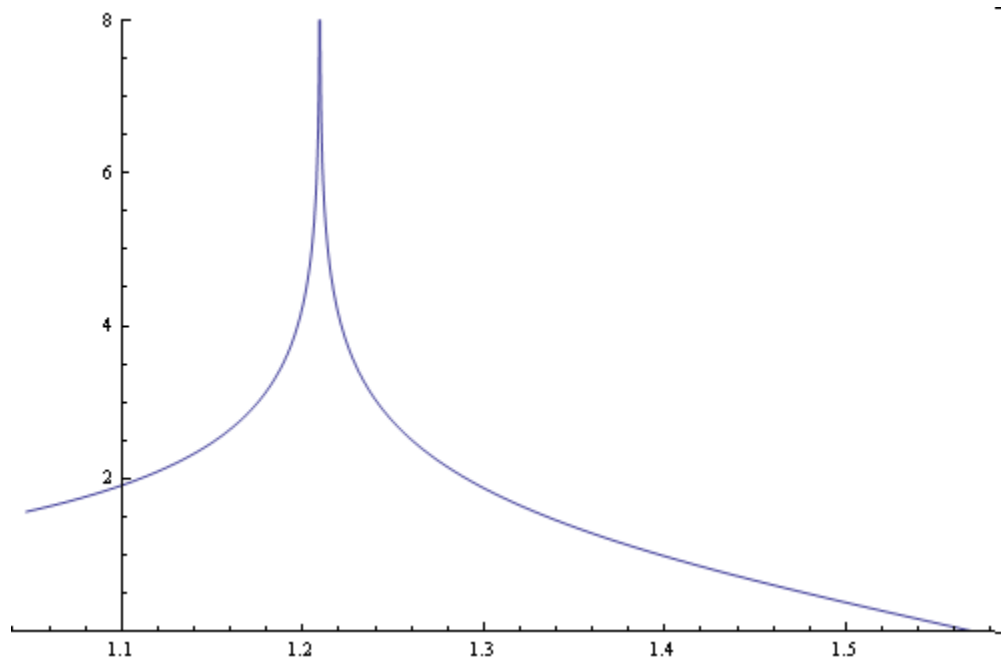

Parallel Performance for the Log-Tan Integral Evaluation

CPUs	Init	Integral #1	Integral #2	Total	Speedup
1	190013	1534652	1026692	2751357	1.00
16	12266	101647	64720	178633	15.40
64	3022	24771	16586	44379	62.00
256	770	6333	4194	11297	243.55
1024	199	1536	1034	2769	993.63

Timings are in seconds. The one-CPU timings are estimated from 16-CPU timings with no communication.

Ising Integrals

We recently applied our methods to study three classes of integrals that arise in the Ising theory of mathematical physics – D_n and two others:

$$C_n := \frac{4}{n!}\int_0^\infty \cdots \int_0^\infty \frac{1}{\left(\sum_{j=1}^n (u_j + 1/u_j)\right)^2} \frac{du_1}{u_1} \cdots \frac{du_n}{u_n}$$

$$D_n := \frac{4}{n!}\int_0^\infty \cdots \int_0^\infty \frac{\prod_{i<j}\left(\frac{u_i - u_j}{u_i + u_j}\right)^2}{\left(\sum_{j=1}^n (u_j + 1/u_j)\right)^2} \frac{du_1}{u_1} \cdots \frac{du_n}{u_n}$$

$$E_n = 2\int_0^1 \cdots \int_0^1 \left(\prod_{1 \le j < k \le n} \frac{u_k - u_j}{u_k + u_j}\right)^2 dt_2\, dt_3 \cdots dt_n$$

where in the last line $u_k = t_1\, t_2 \ldots t_k$.

DHB, J. M. Borwein and R. E. Crandall, "Integrals of the Ising Class," *Journal of Physics A: Mathematical and General*, vol. 39 (2006), pg. 12271-12302.

Computing and Evaluating C_n

We observed that the multi-dimensional C_n integrals can be transformed to 1-D integrals:

$$C_n \quad = \quad \frac{2^n}{n!} \int_0^\infty t K_0^n(t)\, \mathrm{d}t$$

where K_0 is the modified Bessel function. In this form, the C_n constants appear naturally in in quantum field theory (QFT).

We used this formula to compute 1000-digit numerical values of various C_n, from which the following results and others were found, then proven:

$$\begin{aligned}
C_1 &= 2 \\
C_2 &= 1 \\
C_3 &= \mathrm{L}_{-3}(2) = \sum_{n \geq 0} \left(\frac{1}{(3n+1)^2} - \frac{1}{(3n+2)^2} \right) \\
C_4 &= \frac{7}{12} \zeta(3)
\end{aligned}$$

Limiting Value of C_n

The C_n numerical values appear to approach a limit. For instance,

$$C_{1024} = 0.6304735033743867961220401927108789043545870787127323 4\ldots$$

What is this limit? We copied the first 50 digits of this numerical value into the online Inverse Symbolic Calculator (ISC):
http://ddrive.cs.dal.ca/~isc

The result was:

$$\lim_{n\to\infty} C_n \quad = \quad 2e^{-2\gamma}$$

where gamma denotes Euler's constant. Finding this limit led us to the asymptotic expansion and made it clear that the integral representation of C_n is fundamental.

Other Ising Integral Evaluations

$$
\begin{aligned}
D_2 &= 1/3 \\
D_3 &= 8 + 4\pi^2/3 - 27\,\mathrm{L}_{-3}(2) \\
D_4 &= 4\pi^2/9 - 1/6 - 7\zeta(3)/2 \\
E_2 &= 6 - 8\log 2 \\
E_3 &= 10 - 2\pi^2 - 8\log 2 + 32\log^2 2 \\
E_4 &= 22 - 82\zeta(3) - 24\log 2 + 176\log^2 2 - 256(\log^3 2)/3 \\
&\quad +16\pi^2\log 2 - 22\pi^2/3 \\
E_5 &\stackrel{?}{=} 42 - 1984\,\mathrm{Li}_4(1/2) + 189\pi^4/10 - 74\zeta(3) - 1272\zeta(3)\log 2 \\
&\quad +40\pi^2\log^2 2 - 62\pi^2/3 + 40(\pi^2\log 2)/3 + 88\log^4 2 \\
&\quad +464\log^2 2 - 40\log 2
\end{aligned}
$$

where $\mathrm{Li}_n(x)$ is the polylog function. D_2, D_3 and D_4 were originally provided to us by mathematical physicist Craig Tracy, who hoped that our tools could help identify D_5.

The Ising Integral E_5

We were able to reduce E_5, which is a 5-D integral, to an extremely complicated 3-D integral.

We computed this integral to 250-digit precision, using a highly parallel, high-precision 3-D quadrature program. Then we used a PSLQ program to discover the evaluation given on the previous page.

We also computed D_5 to 500 digits, but were unable to identify it. The digits are available if anyone wishes to further explore this question.

$$
\begin{aligned}
E_5 \;=\; & \int_0^1 \int_0^1 \int_0^1 [2(1-x)^2(1-y)^2(1-xy)^2(1-z)^2(1-yz)^2(1-xyz)^2 \\
& (-[4(x+1)(xy+1)\log(2)(y^5z^3x^7 - y^4z^2(4(y+1)z+3)x^6 - y^3z((y^2+1)z^2+4(y+ \\
& 1)z+5)x^5 + y^2(4y(y+1)z^3+3(y^2+1)z^2+4(y+1)z-1)x^4 + y(z(z^2+4z \\
& +5)y^2+4(z^2+1)y+5z+4)x^3 + ((-3z^2-4z+1)y^2-4zy+1)x^2 - (y(5z+4) \\
& +4)x-1)] / [(x-1)^3(xy-1)^3(xyz-1)^3] + [3(y-1)^2y^4(z-1)^2z^2(yz \\
& -1)^2x^6 + 2y^3z(3(z-1)^2z^3y^5 + z^2(5z^3+3z^2+3z+5)y^4 + (z-1)^2z \\
& (5z^2+16z+5)y^3 + (3z^5+3z^4-22z^3-22z^2+3z+3)y^2 + 3(-2z^4+z^3+2 \\
& z^2+z-2)y + 3z^3+5z^2+5z+3)x^5 + y^2(7(z-1)^2z^4y^6 - 2z^3(z^3+15z^2 \\
& +15z+1)y^5 + 2z^2(-21z^4+6z^3+14z^2+6z-21)y^4 - 2z(z^5-6z^4-27z^3 \\
& -27z^2-6z+1)y^3 + (7z^6-30z^5+28z^4+54z^3+28z^2-30z+7)y^2 - 2(7z^5 \\
& +15z^4-6z^3-6z^2+15z+7)y + 7z^4-2z^3-42z^2-2z+7)x^4 - 2y(z^3(z^3 \\
& -9z^2-9z+1)y^6 + z^2(7z^4-14z^3-18z^2-14z+7)y^5 + z(7z^5+14z^4+3 \\
& z^3+3z^2+14z+7)y^4 + (z^6-14z^5+3z^4+84z^3+3z^2-14z+1)y^3 - 3(3z^5 \\
& +6z^4-z^3-z^2+6z+3)y^2 - (9z^4+14z^3-14z^2+14z+9)y + z^3+7z^2+7z \\
& +1)x^3 + (z^2(11z^4+6z^3-66z^2+6z+11)y^6 + 2z(5z^5+13z^4-2z^3-2z^2 \\
& +13z+5)y^5 + (11z^6+26z^5+44z^4-66z^3+44z^2+26z+11)y^4 + (6z^5-4 \\
& z^4-66z^3-66z^2-4z+6)y^3 - 2(33z^4+2z^3-22z^2+2z+33)y^2 + (6z^3+26 \\
& z^2+26z+6)y + 11z^2+10z+11)x^2 - 2(z^2(5z^3+3z^2+3z+5)y^5 + z(22z^4 \\
& +5z^3-22z^2+5z+22)y^4 + (5z^5+5z^4-26z^3-26z^2+5z+5)y^3 + (3z^4- \\
& 22z^3-26z^2-22z+3)y^2 + (3z^3+5z^2+5z+3)y + 5z^2+22z+5)x + 15z^2+2z \\
& +2y(z-1)^2(z+1) + 2y^3(z-1)^2z(z+1) + y^4z^2(15z^2+2z+15) + y^2(15z^4 \\
& -2z^3-90z^2-2z+15) + 15] / [(x-1)^2(y-1)^2(xy-1)^2(z-1)^2(yz-1)^2 \\
& (xyz-1)^2] - [4(x+1)(y+1)(yz+1)(-z^2y^4+4z(z+1)y^3+(z^2+1)y^2 \\
& -4(z+1)y + 4x(y^2-1)(y^2z^2-1) + x^2(z^2y^4-4z(z+1)y^3-(z^2+1)y^2 \\
& +4(z+1)y+1) - 1)\log(x+1)] / [(x-1)^3x(y-1)^3(yz-1)^3] - [4(y+1)(xy \\
& +1)(z+1)(x^2(z^2-4z-1)y^4 + 4x(x+1)(z^2-1)y^3 - (x^2+1)(z^2-4z-1) \\
& y^2 - 4(x+1)(z^2-1)y + z^2-4z-1)\log(xy+1)] / [x(y-1)^3y(xy-1)^3(z- \\
& 1)^3] - [4(z+1)(yz+1)(x^3y^5z^7 + x^2y^4(4x(y+1)+5)z^6 - xy^3((y^2+ \\
& 1)x^2-4(y+1)x-3)z^5 - y^2(4y(y+1)x^3+5(y^2+1)x^2+4(y+1)x+1)z^4+ \\
& y(y^2x^3-4y(y+1)x^2-3(y^2+1)x-4(y+1))z^3 + (5x^2y^2+y^2+4x(y+1) \\
& y+1)z^2 + ((3x+4)y+4)z-1)\log(xyz+1)] / [xy(z-1)^3z(yz-1)^3(xyz-1)^3])] \\
& / [(x+1)^2(y+1)^2(xy+1)^2(z+1)^2(yz+1)^2(xyz+1)^2] \; dx\, dy\, dz
\end{aligned}
$$

Recursions in Ising Integrals

Consider the 2-parameter class of Ising integrals (which arises in QFT for odd *k*):

$$C_{n,k} = \frac{4}{n!} \int_0^\infty \cdots \int_0^\infty \frac{1}{\left(\sum_{j=1}^n (u_j + 1/u_j)\right)^{k+1}} \frac{\mathrm{d}u_1}{u_1} \cdots \frac{\mathrm{d}u_n}{u_n}$$

After computing 1000-digit numerical values for all *n* up to 36 and all *k* up to 75 (performed on a highly parallel computer system), we discovered (using PSLQ) linear relations in the rows of this array. For example, when *n* = 3:

$$\begin{aligned}
0 &= C_{3,0} - 84C_{3,2} + 216C_{3,4} \\
0 &= 2C_{3,1} - 69C_{3,3} + 135C_{3,5} \\
0 &= C_{3,2} - 24C_{3,4} + 40C_{3,6} \\
0 &= 32C_{3,3} - 630C_{3,5} + 945C_{3,7} \\
0 &= 125C_{3,4} - 2172C_{3,6} + 3024C_{3,8}
\end{aligned}$$

Similar, but more complicated, recursions have been found for all *n*.

DHB, D. Borwein, J. M. Borwein and R. Crandall, “Hypergeometric Forms for Ising-Class Integrals,” *Experimental Mathematics*, to appear, http://crd.lbl.gov/~dhbailey/dhbpapers/meijer/pdf.

J. M. Borwein and B. Salvy, “A Proof of a Recursion for Bessel Moments,” *Experimental Mathematics*, vol. 17 (2008), pg. 223-230.

Four Hypergeometric Evaluations

$$
\begin{aligned}
c_{3,0} &= \frac{3\Gamma^6(1/3)}{32\pi 2^{2/3}} = \frac{\sqrt{3}\pi^3}{8} {}_3F_2\left(\begin{array}{c} 1/2, 1/2, 1/2 \\ 1, 1 \end{array} \middle| \frac{1}{4}\right) \\
c_{3,2} &= \frac{\sqrt{3}\pi^3}{288} {}_3F_2\left(\begin{array}{c} 1/2, 1/2, 1/2 \\ 2, 2 \end{array} \middle| \frac{1}{4}\right) \\
c_{4,0} &= \frac{\pi^4}{4} \sum_{n=0}^{\infty} \frac{\binom{2n}{n}^4}{4^{4n}} = \frac{\pi^4}{4} {}_4F_3\left(\begin{array}{c} 1/2, 1/2, 1/2, 1/2 \\ 1, 1, 1 \end{array} \middle| 1\right) \\
c_{4,2} &= \frac{\pi^4}{64} \left[4 {}_4F_3\left(\begin{array}{c} 1/2, 1/2, 1/2, 1/2 \\ 1, 1, 1 \end{array} \middle| 1\right)\right. \\
& \qquad \left. -3 {}_4F_3\left(\begin{array}{c} 1/2, 1/2, 1/2, 1/2 \\ 2, 1, 1 \end{array} \middle| 1\right)\right] - \frac{3\pi^2}{16}
\end{aligned}
$$

DHB, J. M. Borwein, D. Broadhurst and M. L. Glasser, “Elliptic Integral Evaluations of Bessel Moments,” *Journal of Physics A: Mathematical and General*, vol. 41 (2008), pg 205203.

2-D Integral in Bessel Moment Study

We conjectured (and later proved)

$$c_{5,0} \quad = \quad \frac{\pi}{2}\int_{-\pi/2}^{\pi/2}\int_{-\pi/2}^{\pi/2}\frac{\mathbf{K}(\sin\theta)\,\mathbf{K}(\sin\phi)}{\sqrt{\cos^2\theta\cos^2\phi+4\sin^2(\theta+\phi)}}\,\mathrm{d}\theta\,\mathrm{d}\phi$$

Here **K** denotes the complete elliptic integral of the first kind

Note that the integrand function has singularities on all four sides of the region of integration.

We were able to evaluate this integral to 120-digit accuracy, using 1024 cores of the "Franklin" Cray XT4 system at LBNL.

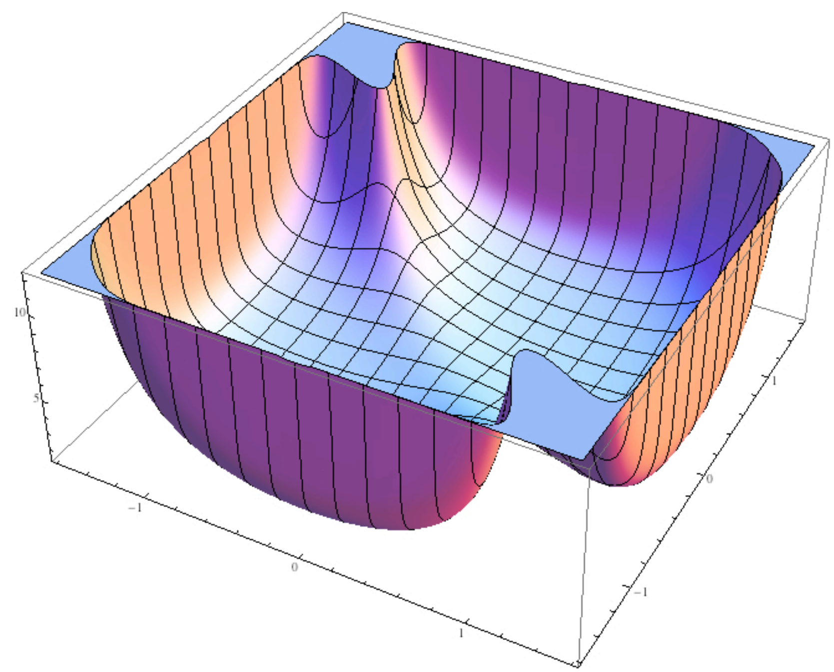

Heisenberg Spin Integrals

In another recent application of these methods, we investigated the following "spin integrals," which arise from studies in mathematical physics:

$$P(n) \quad := \quad \frac{\pi^{n(n+1)/2}}{(2\pi i)^n} \cdot \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \cdots \int_{-\infty}^{\infty} U(x_1 - i/2, x_2 - i/2, \cdots, x_n - i/2)$$
$$\times\, T(x_1 - i/2, x_2 - i/2, \cdots, x_n - i/2)\, \mathrm{d}x_1 \mathrm{d}x_2 \cdots \mathrm{d}x_n$$

where

$$U(x_1 - i/2, x_2 - i/2, \cdots, x_n - i/2) \;=\; \frac{\prod_{1 \le k < j \le n} \sinh[\pi(x_j - x_k)]}{\prod_{1 \le j \le n} i^n \cosh^n(\pi x_j)}$$

$$T(x_1 - i/2, x_2 - i/2, \cdots, x_n - i/2) \;=\; \frac{\prod_{1 \le j \le n} (x_j - i/2)^{j-1} (x_j + i/2)^{n-j}}{\prod_{1 \le k < j \le n} (x_j - x_k - i)}$$

H. E. Boos, V. E. Korepin, Y. Nishiyama and M. Shiroishi, "Quantum Correlations and Number Theory," *Journal of Physics A: Mathematical and General*, vol. 35 (2002), pg. 4443.

Evaluations of *P*(*n*) Derived Analytically, Confirmed Numerically

$$
\begin{aligned}
P(1) &= \frac{1}{2}, \quad P(2) = \frac{1}{3} - \frac{1}{3}\log 2, \quad P(3) = \frac{1}{4} - \log 2 + \frac{3}{8}\zeta(3) \\
P(4) &= \frac{1}{5} - 2\log 2 + \frac{173}{60}\zeta(3) - \frac{11}{6}\zeta(3)\log 2 - \frac{51}{80}\zeta^2(3) - \frac{55}{24}\zeta(5) + \frac{85}{24}\zeta(5)\log 2 \\
P(5) &= \frac{1}{6} - \frac{10}{3}\log 2 + \frac{281}{24}\zeta(3) - \frac{45}{2}\zeta(3)\log 2 - \frac{489}{16}\zeta^2(3) - \frac{6775}{192}\zeta(5) \\
&\quad + \frac{1225}{6}\zeta(5)\log 2 - \frac{425}{64}\zeta(3)\zeta(5) - \frac{12125}{256}\zeta^2(5) + \frac{6223}{256}\zeta(7) \\
&\quad - \frac{11515}{64}\zeta(7)\log 2 + \frac{42777}{512}\zeta(3)\zeta(7)
\end{aligned}
$$

and a much more complicated expression for *P*(6). Run times increase very rapidly with the dimension *n*:

n	Digits	Processors	Run Time
2	120	1	10 sec.
3	120	8	55 min.
4	60	64	27 min.
5	30	256	39 min.
6	6	256	59 hrs.

Box Integrals

The following integrals appear in studies, say, of the average distance between points in a cube, or the average electric potential in a cube:

$$B_n(s) \quad := \quad \int_0^1 \cdots \int_0^1 \left(r_1^2 + \cdots + r_n^2\right)^{s/2} \; dr_1 \cdots dr_n$$

$$\Delta_n(s) \quad := \quad \int_0^1 \cdots \int_0^1 \left((r_1 - q_1)^2 + \cdots + (r_n - q_n)^2\right)^{s/2} \; dr_1 \cdots dr_n \; dq_1 \cdots dq_n$$

DHB, J. M. Borwein and R. E. Crandall, “Box Integrals,” *Journal of Computational and Applied Mathematics*, vol. 206 (2007), pg. 196-208.

Evaluations of Box Integrals

$$
\begin{aligned}
B_2(-1) &= \log\left(3+2\sqrt{2}\right) \\
B_3(-1) &= -\frac{\pi}{4} - \frac{1}{2}\log 2 + \log(5+3\sqrt{3}) \\
B_1(1) &= \frac{1}{2} \\
B_2(1) &= \frac{\sqrt{2}}{3} + \frac{1}{3}\log\left(\sqrt{2}+1\right) \\
B_3(1) &= \frac{\sqrt{3}}{4} + \frac{1}{2}\log\left(2+\sqrt{3}\right) - \frac{\pi}{24} \\
B_4(1) &= \frac{2}{5} + \frac{7}{20}\pi\sqrt{2} - \frac{1}{20}\pi\log\left(1+\sqrt{2}\right) + \log(3) - \frac{7}{5}\sqrt{2}\arctan\left(\sqrt{2}\right) + \frac{1}{10}\mathcal{K}_0
\end{aligned}
$$

where

$$
\mathcal{K}_0 := \int_0^1 \frac{\log(1+\sqrt{3+y^2}) - \log(-1+\sqrt{3+y^2})}{1+y^2}\,dy = 2\int_0^1 \frac{\operatorname{arctanh}\left(\frac{1}{\sqrt{3+y^2}}\right)}{1+y^2}\,dy
$$

Evaluations of Box Integrals, Cont.

$$
\begin{aligned}
\Delta_2(-1) &= \frac{4}{3} - \frac{4}{3}\sqrt{2} + 4\log(1+\sqrt{2}) \\
\Delta_1(1) &= \frac{1}{3} \\
\Delta_2(1) &= \frac{1}{15}\left(2 + \sqrt{2} + 5\log(1+\sqrt{2})\right), \\
\Delta_3(1) &= \frac{4}{105} + \frac{17}{105}\sqrt{2} - \frac{2}{35}\sqrt{3} + \frac{1}{5}\log(1+\sqrt{2}) + \frac{2}{5}\log(2+\sqrt{3}) - \frac{1}{15}\pi, \\
\Delta_4(1) &= \frac{26}{15}\,\mathrm{G} - \frac{34}{105}\,\pi\sqrt{2} - \frac{16}{315}\,\pi + \frac{197}{420}\,\log(3) + \frac{52}{105}\,\log\left(2+\sqrt{3}\right) \\
&+ \frac{1}{14}\,\log\left(1+\sqrt{2}\right) + \frac{8}{105}\,\sqrt{3} + \frac{73}{630}\,\sqrt{2} - \frac{23}{135} + \frac{136}{105}\,\sqrt{2}\arctan\left(\frac{1}{\sqrt{2}}\right) \\
&- \frac{1}{5}\,\pi\,\log\left(1+\sqrt{2}\right) + \frac{4}{5}\,\alpha\,\log\left(1+\sqrt{2}\right) - \frac{4}{5}\,\mathrm{Cl}_2\left(\alpha\right) - \frac{4}{5}\,\mathrm{Cl}_2\left(\alpha + \frac{\pi}{2}\right)
\end{aligned}
$$

where G is Catalan's constant and Cl denotes the Clausen function.

New Result (18 Jan 2009)

$$\begin{aligned}\Delta_3(-1) &= \frac{2}{\sqrt{\pi}} \int_0^\infty \frac{(-1 + e^{-u^2} + \sqrt{\pi}\, u \operatorname{erf}(u))^3}{u^6} \, \mathrm{d}u \\ &= \frac{1}{15} \left(6 + 6\sqrt{2} - 12\sqrt{3} - 10\pi + 30 \log(1 + \sqrt{2}) + 30 \log(2 + \sqrt{3})\right)\end{aligned}$$

As in many of the previous results, this was found by first computing the integral to high precision (1000 digits in this case), conjecturing possible terms on the right-hand side, then applying PSLQ to look for a relation. We now have proven this result.

This and similar integrals have recently arisen in problems suggested by neuroscientists – e.g., the average distance between synapses in a mouse brain.

Ref: Work in progress! Will be written up soon.

Summary

- Achieving high fractions of peak performance on emerging petascale, multicore-based systems will be a major challenge.
- Semi-automatic tuning techniques may be effective in improving local node performance on multicore processor nodes, but these will not be widely available for several years.
- One bright spot is the emerging need for high-precision arithmetic, which maps well onto state-of-the-art computer architectures.
- The emerging “experimental” methodology in mathematics and mathematical physics requires hundreds or thousands of digits.
- High-precision evaluation of integrals, followed by constant-recognition techniques, has been a particularly fruitful area of recent research.