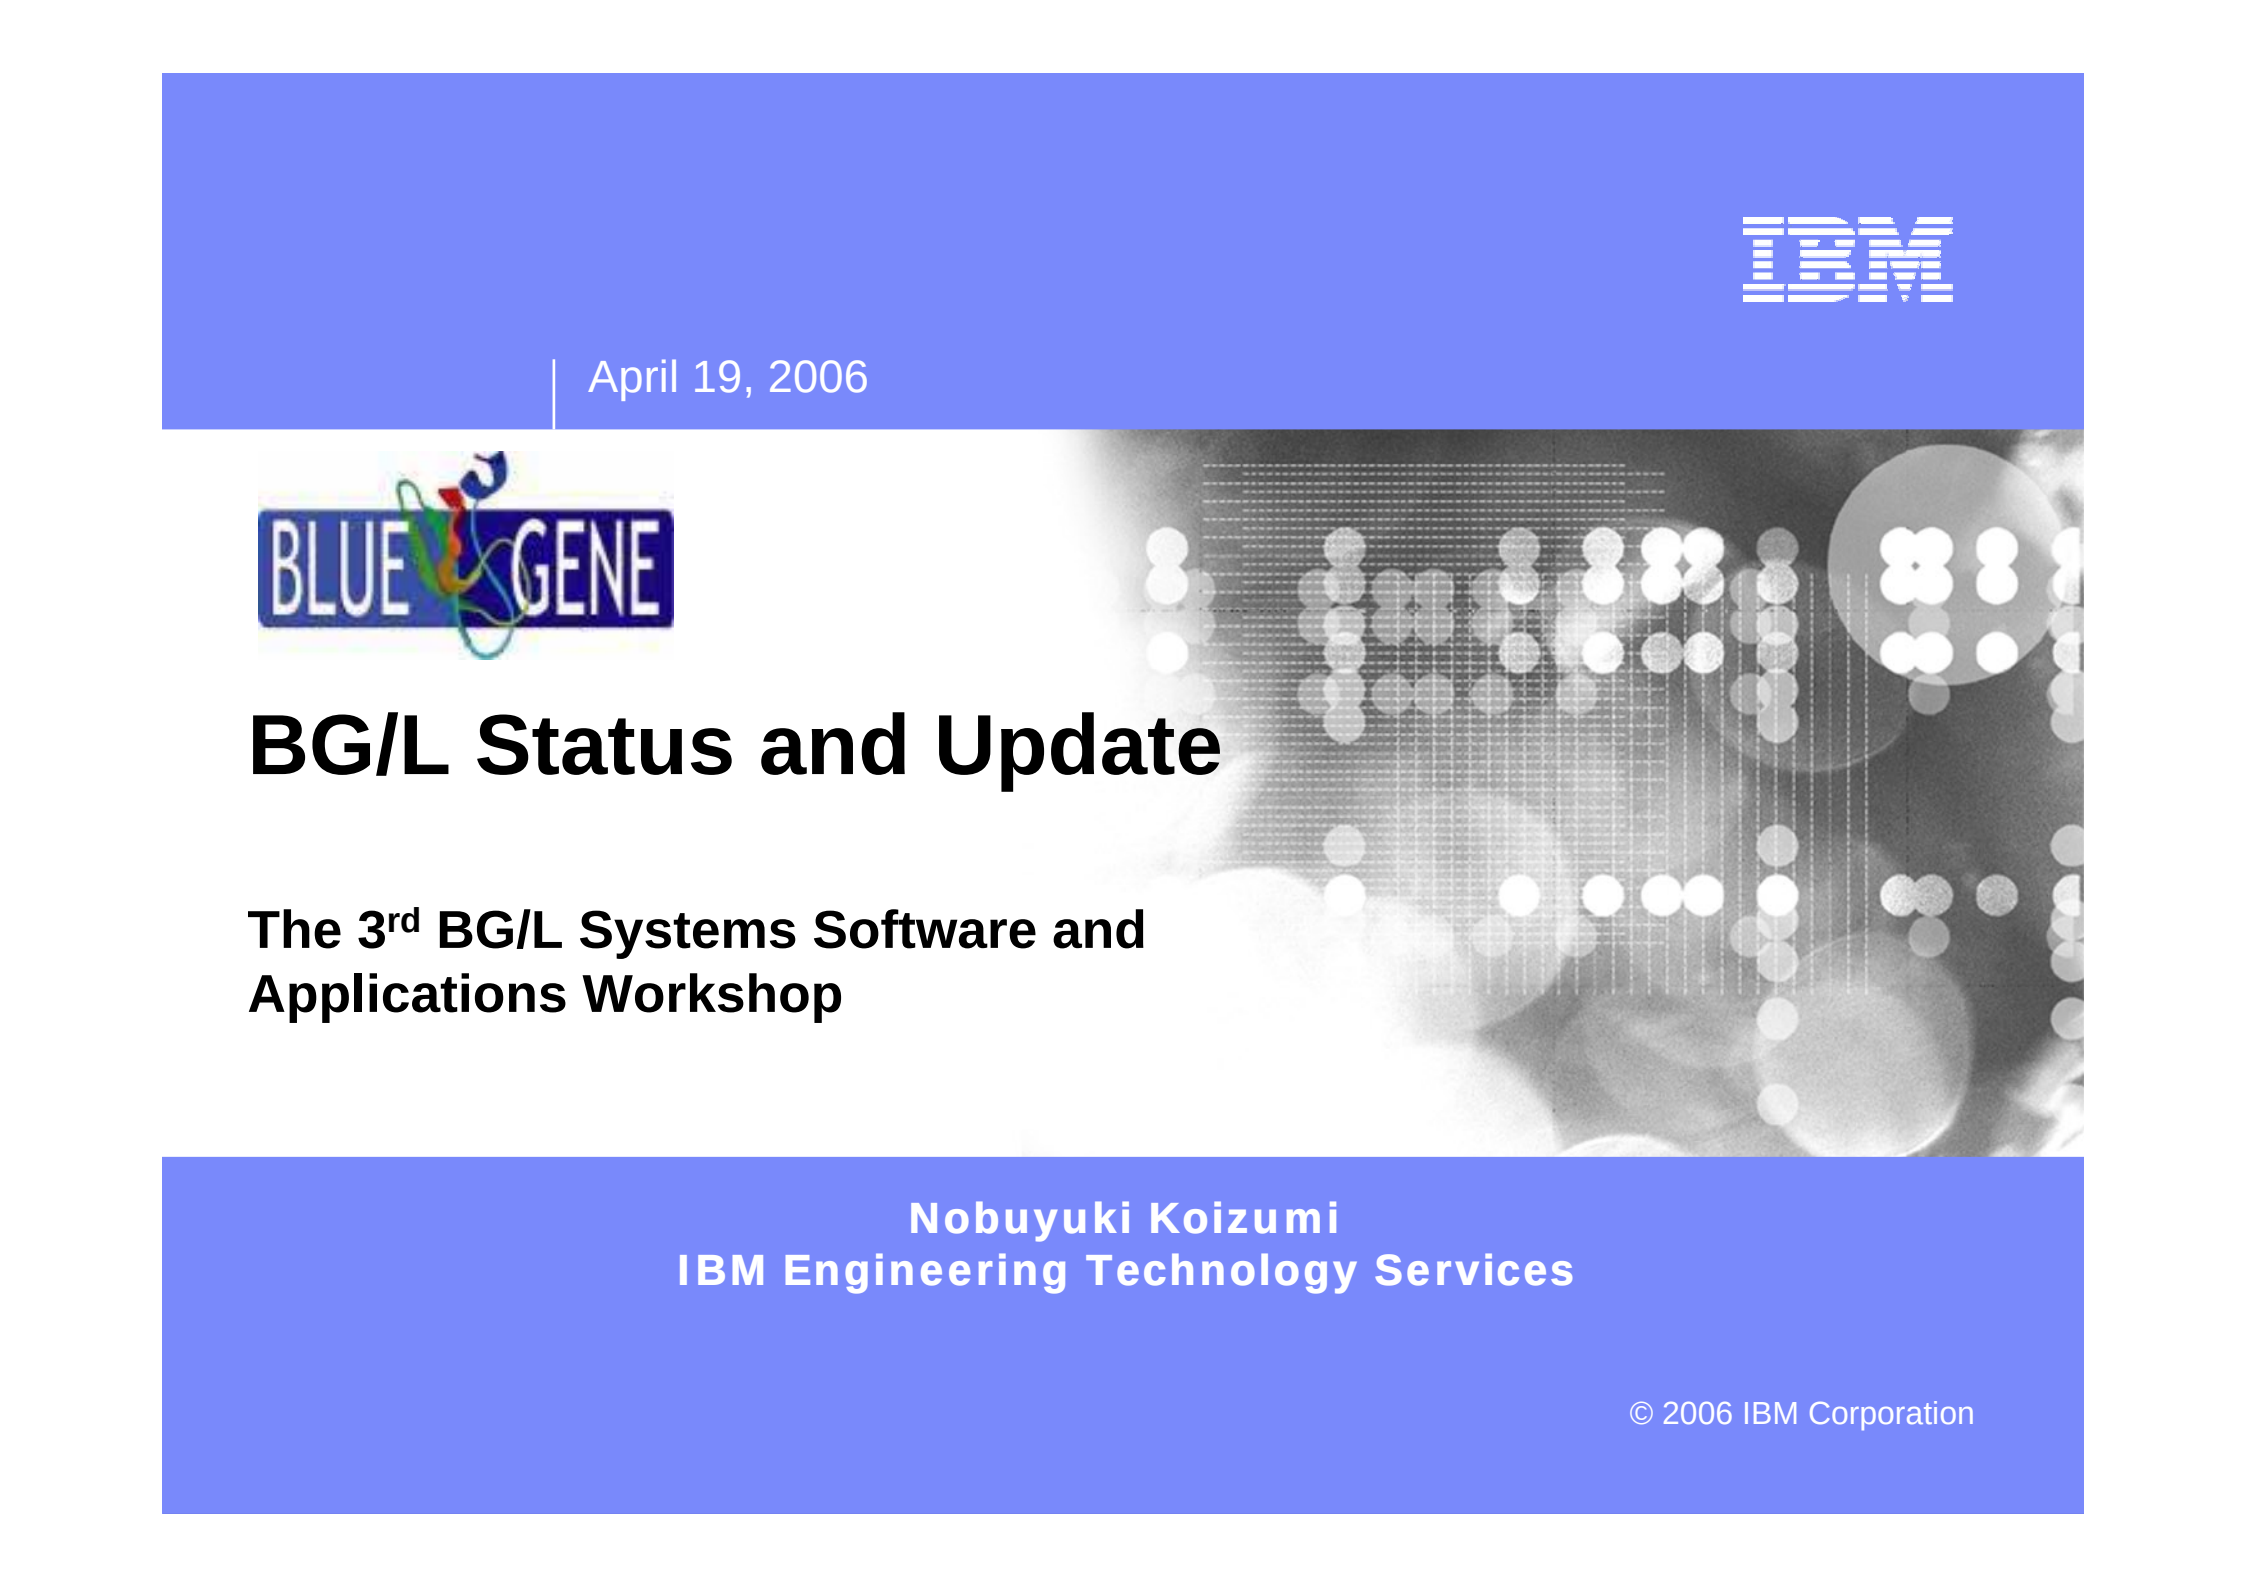

BG/L Status and Update

The 3rd BG/L Systems Software and Applications Workshop

Nobuyuki Koizumi
IBM Engineering Technology Services

What is driving the need for more compute cycles?

Science Area	Breakthrough Target Requiring 100's of Teraflops
Fusion	Simulate hydrogen plasma torus (ITER reactor) operating at over 100 million ℃ to produce 500 MW of fusion power
Climate	Perform full ocean / atmosphere climate model with 0.125 degree spacing and an ensemble of 8-10 runs
Nanoscience	Simulate nanostructures with 100's to 1000's of atoms
Combustion	Simulate laboratory scale flames with high fidelity
Astrophysics	Simulate explosion of a supernova with a full 3D model

Source: "Towards Petascale Computing for Science," Horst Simon, Lawrence Berkeley National Laboratory, ICCSE 2005 Istanbul, June 2005

What is driving the need for more compute cycles?

Commercial Area	Breakthrough Target Requiring 100's of TeraFLOP/S
Petroleum	Run PSTM and PSDM routines for large-scale seismic surveys in < ½ day
Finance	Perform Monte Carlo simulations using > 500M variables for risk analysis, options pricing, derivatives hedging
Life Sciences	Conduct real-time simulations of the brain's neocortical column using a cellular-level model
Engineering	Perform CFD simulation of full airframe using 100's more mesh points than required to simulate an airplane wing
Supply Chain Management	Model optimized deployment of 1000's of part numbers across 100's of parts depots with turnaround time < 1 hr

What is the roadblock to more compute cycles?

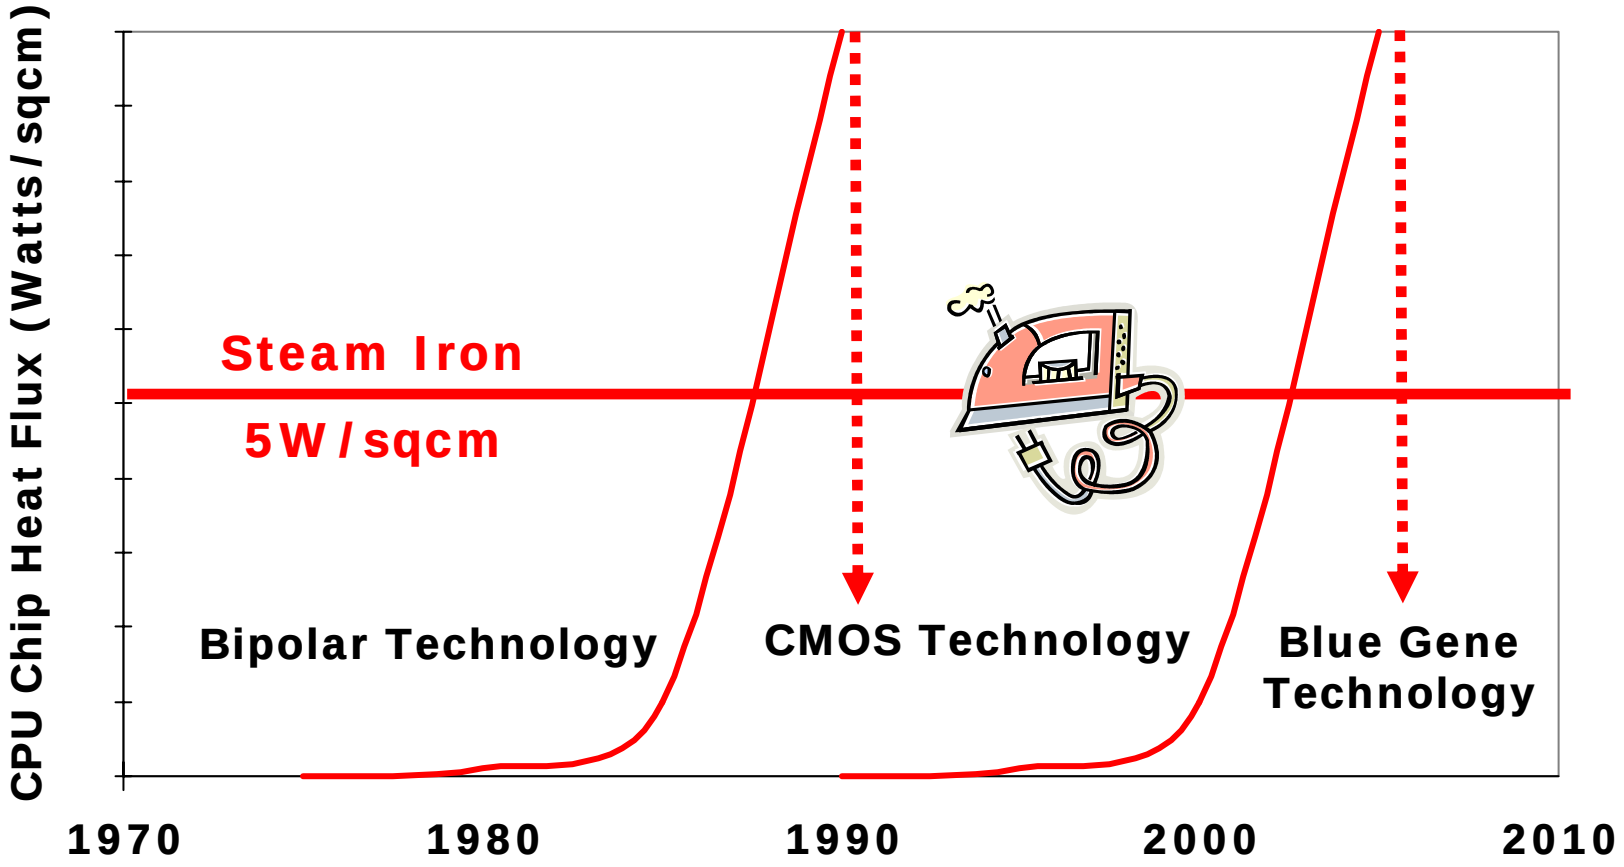

Blue Gene is an evolutionary, innovative technology which reduces "time to solution" for many computational science problems through ultrascalability and modularity with the lowest power consumption, smallest footprint, highest reliability and easiest manageability in the industry.

What is Blue Gene?

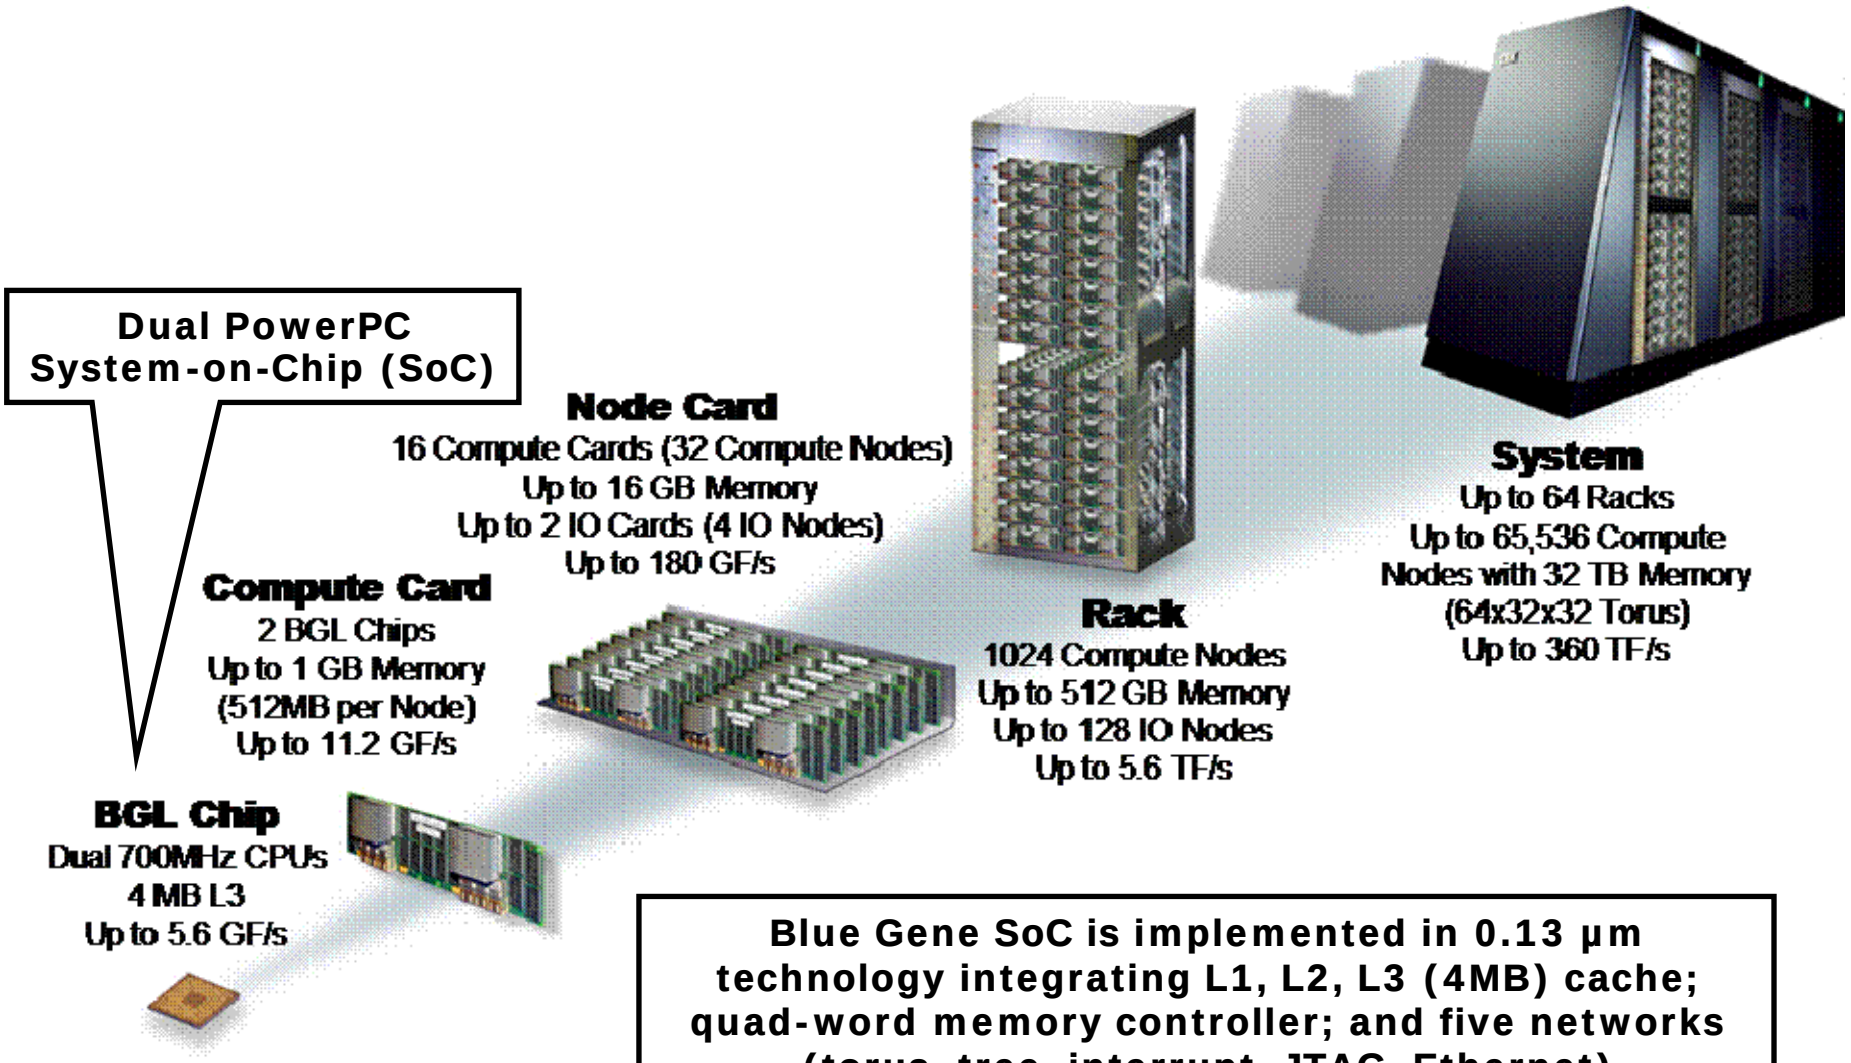

Blue Gene SoC is implemented in 0.13 µm technology integrating L1, L2, L3 (4MB) cache; quad-word memory controller; and five networks (torus, tree, interrupt, JTAG, Ethernet)

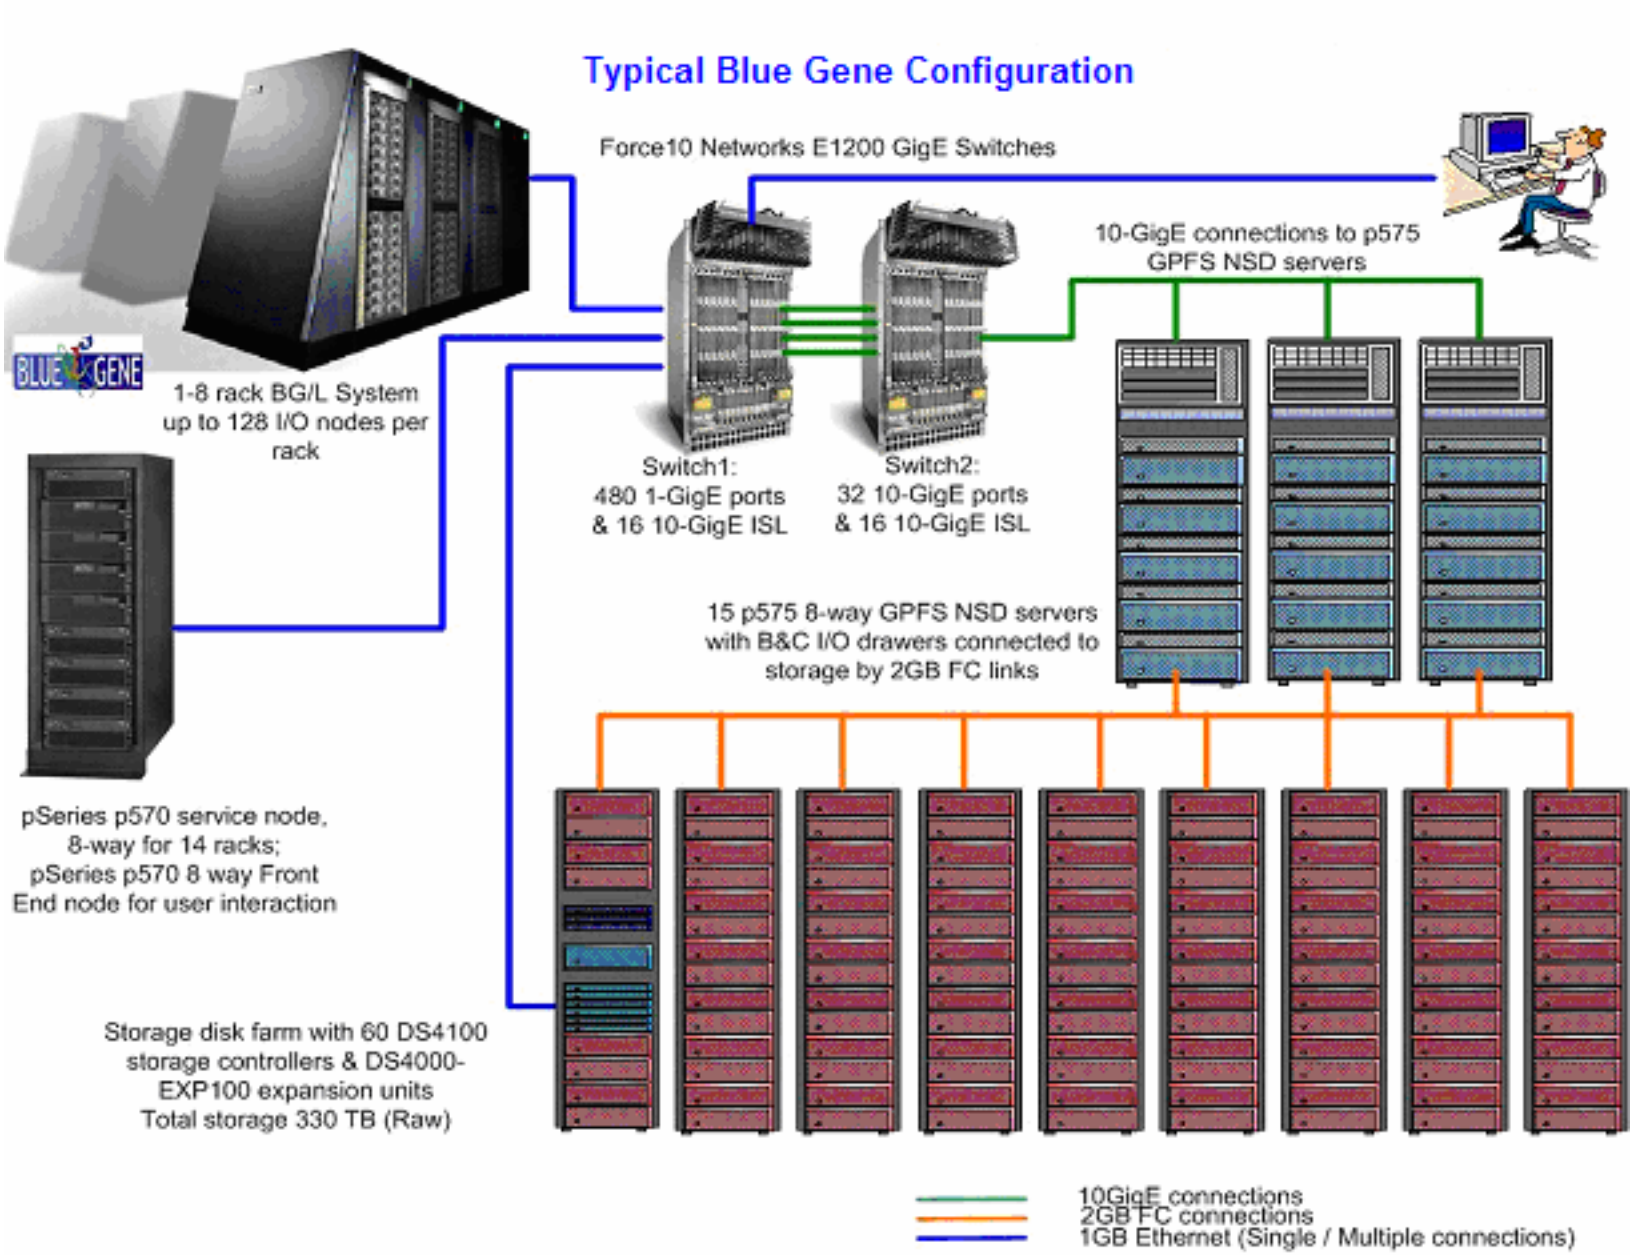

Blue Gene balances massive scale-out capacity while preserving familiar user/administrator environments

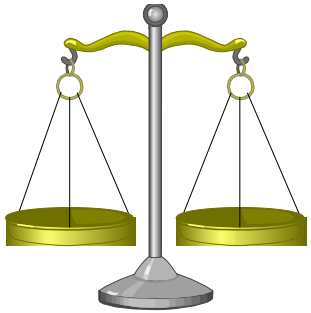

- Scalability to > 65K processors
- Transparent high-speed, low latency networking
- Standards-based Message Passing Interface (MPI)
- Familiar, standard programming environment
  - ✓ Linux based development environment (SuSE SLES9)
  - ✓ Blue Gene's Compute Node Kernel (CNK) provides POSIX system calls with restrictions to ensure scaling
  - ✓ Automatic SIMD (Single-Instruction Multiple-Data) FPU exploitation enabled by Fortran, C, and C++ compilers

Energy and space savings with Blue Gene

"Commodity microprocessors and Linux cluster architectures offer the potential for unprecedented levels of hardware price/performance. They also, however, dramatically increase data center space, power and cooling requirements. Costs for these may more than offset the advantages of inexpensive hardware, and may create reliability and manageability problems that impair system effectiveness. It is from this perspective that the potential role of the IBM Blue Gene system should be viewed."

Business Case for IBM eServer Blue Gene Systems: Reducing the Costs of Commercial Supercomputing, International Technology Group, Nov. 2005

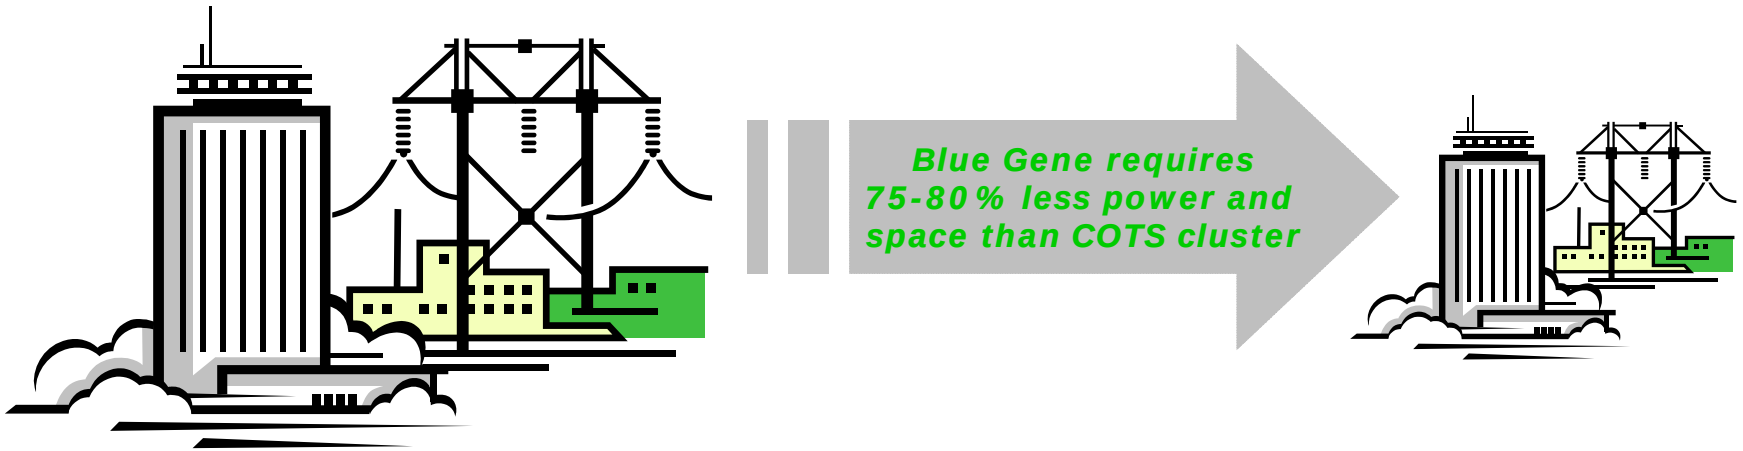

Comparative 3-year infrastructure costs *(existing facilities)*

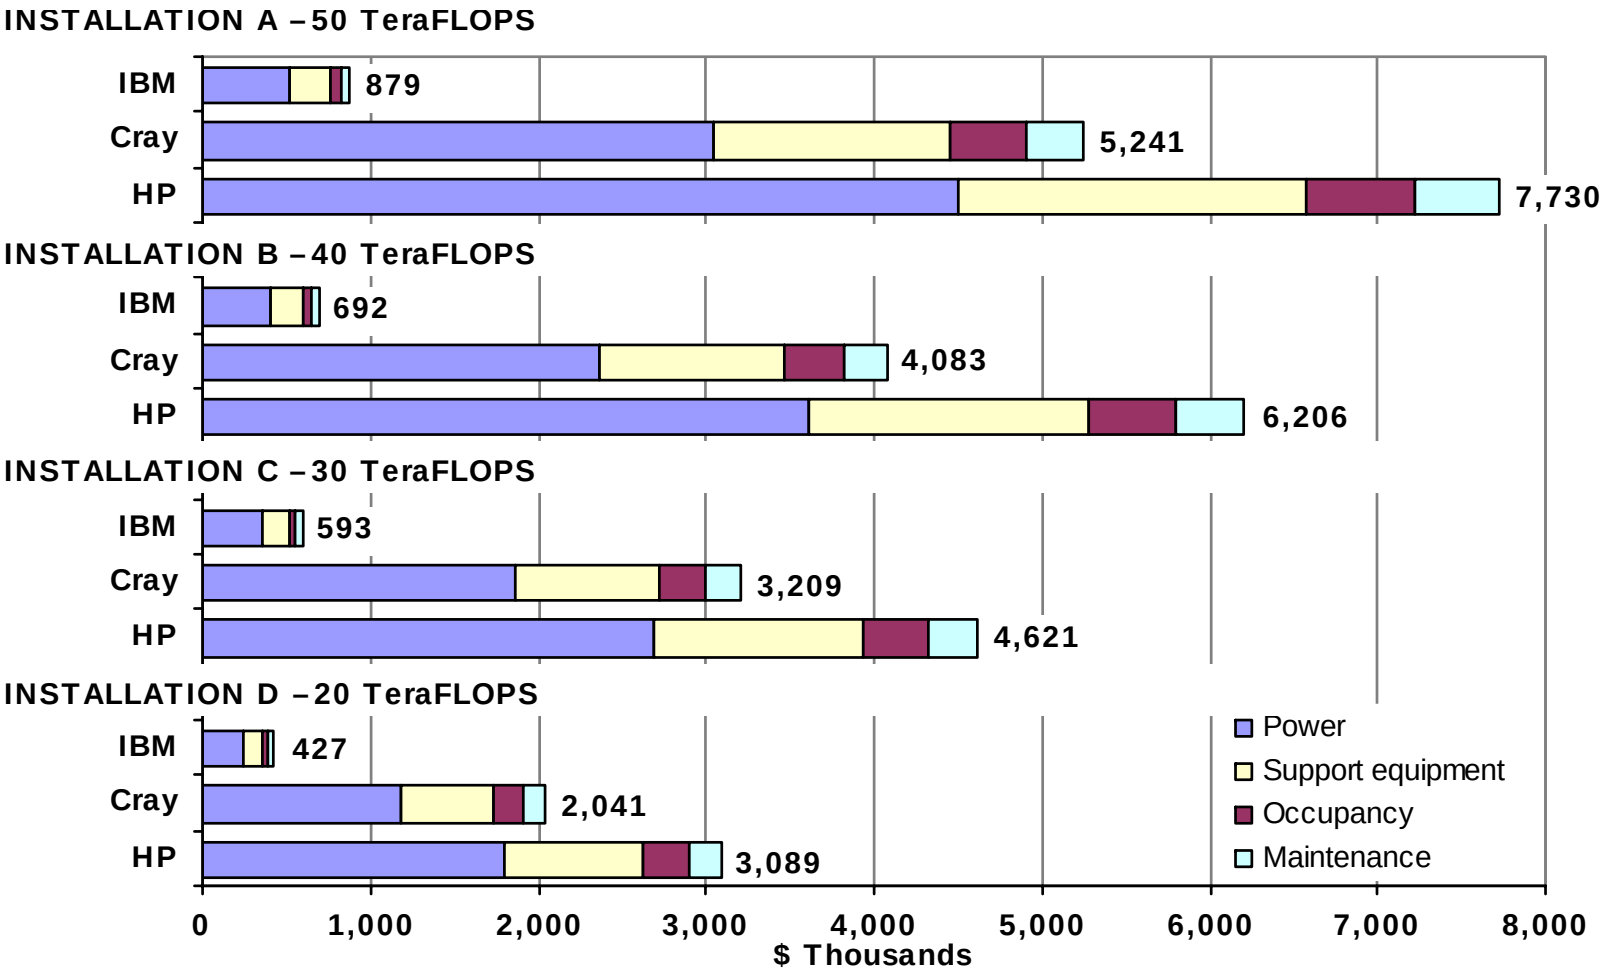

Business Case for IBM eServer Blue Gene Systems: Reducing the Costs of Commercial Supercomputing, International Technology Group, November 2005

Blue Gene power and space efficiency

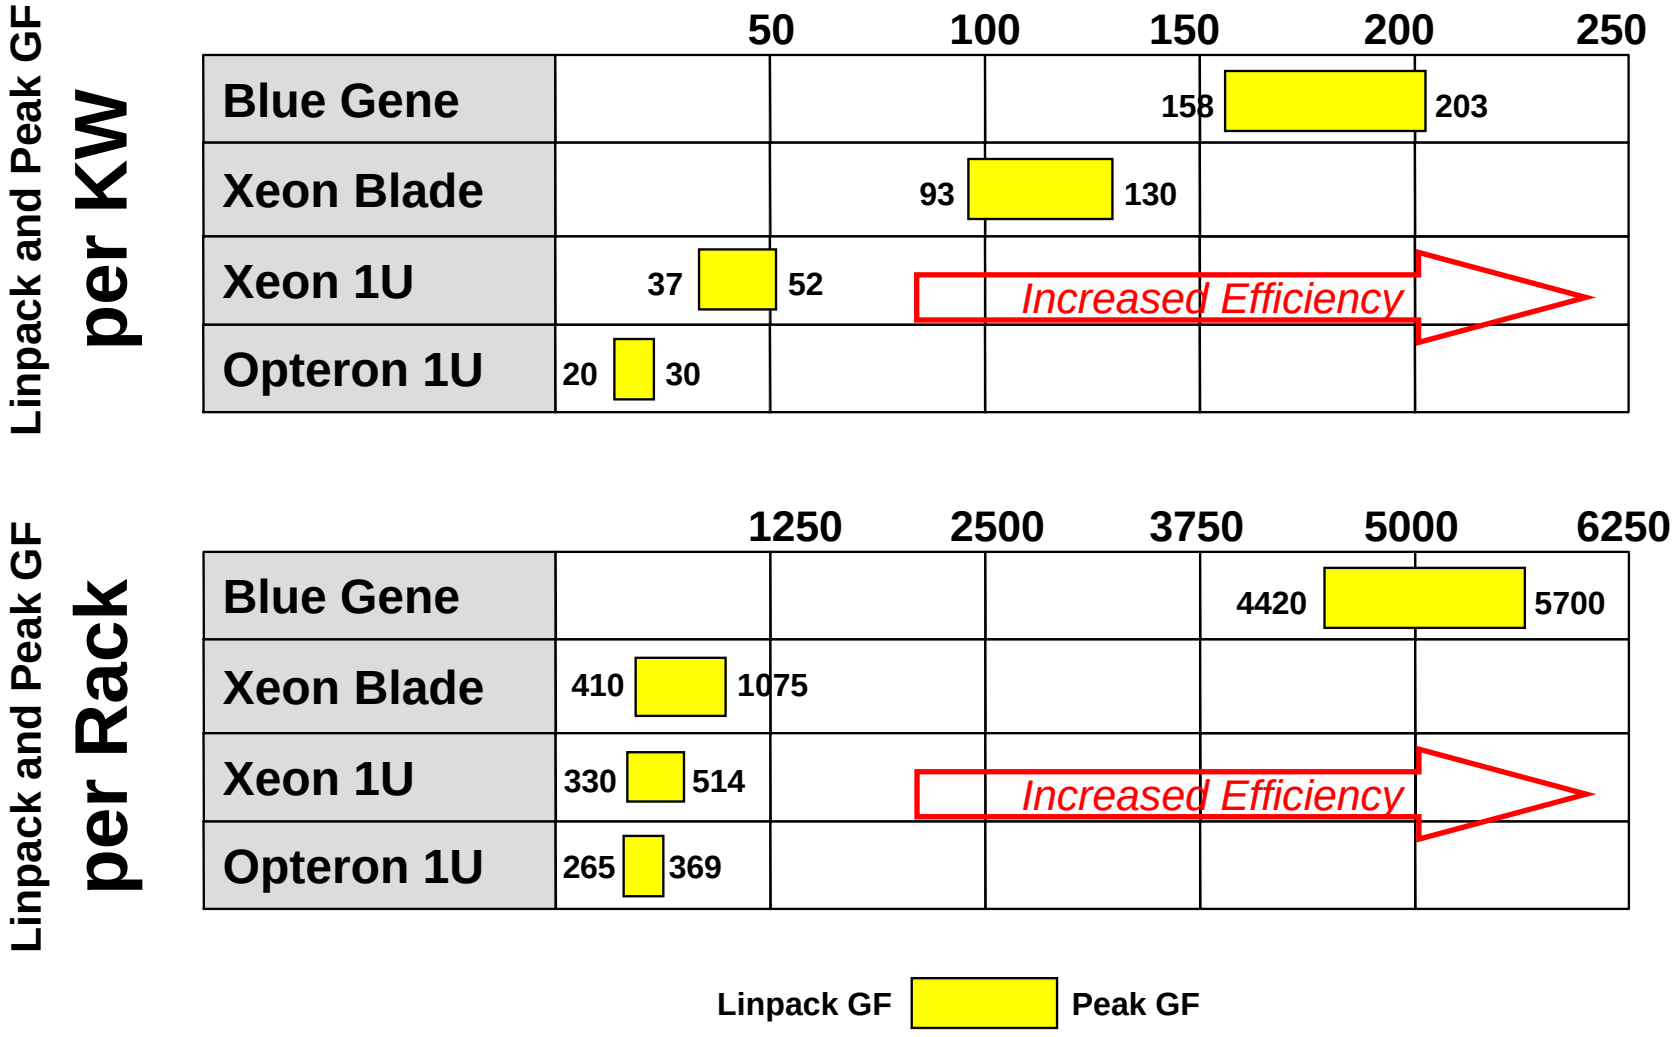

Blue Gene award-winning performance

- 3 in Top10 (#1 and #2)
- 7 in Top50 (11-50)
- 9 in Top100 (51-100)
- 19 in Top500

www.top500.org

HPC Challenge Benchmarks	IBM Blue Gene 65,536 nodes 64 racks	Cray XT3 (Opteron) 5200 nodes 56 racks
HPL (TFLOP/s) – Linpack TPP benchmark which measures floating point rate of execution for solving linear system of equations	259.21	20.53
RANDOMACCESS (GUP/s) – measures rate of integer random updates of memory	35.46	0.69 (7.69 on Cray X1E)
FETE (GFLOP/s) – measures floating point rate of execution of double precision complex one-dimensional Discrete Fourier Transform (DFT)	2311.09	905.57
STREAM (GB/s) – simple synthetic that measures sustainable memory bandwidth and corresponding computation rate for simple vector kernel	160,064	26,021

Where in the world is Blue Gene?

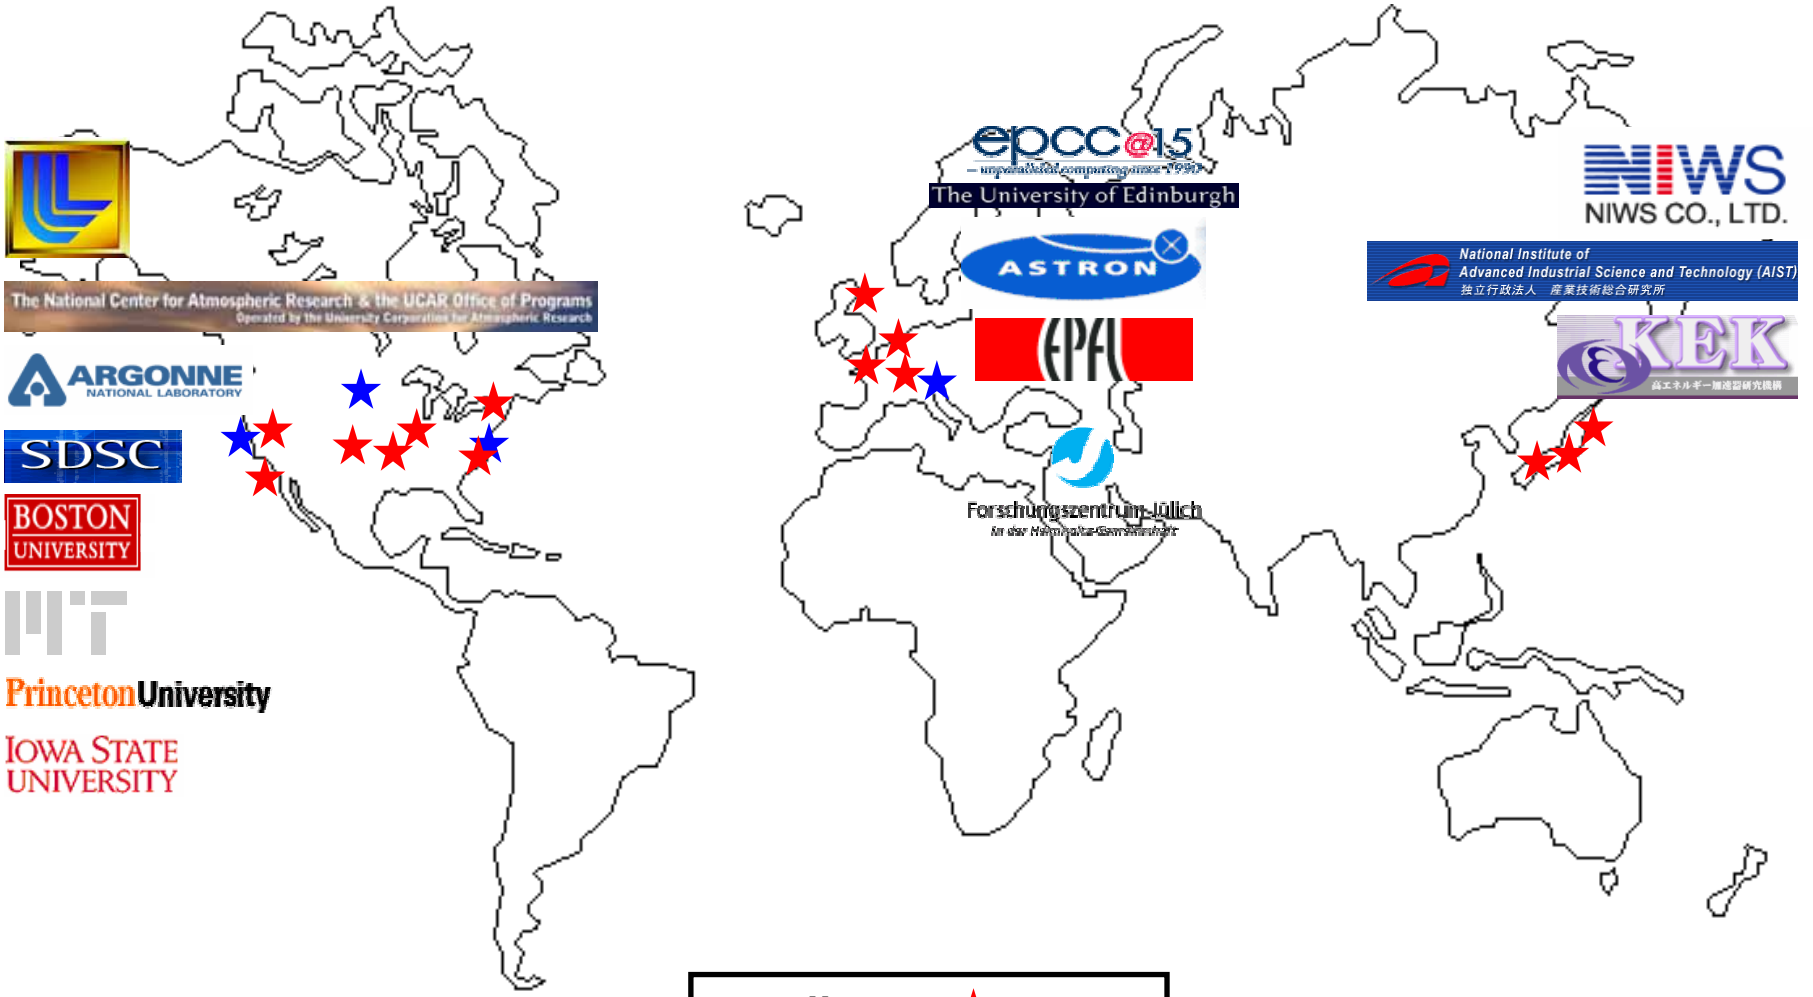

How are they using Blue Gene?

Weapons Research

Lawrence Livermore National Laboratory in California has the cream of the Blue Gene crop: a 64-rack machine … used for nuclear-weapons research.

Radioastronomy

The Netherlands' Low Frequency Array radio telescope uses a six-rack Blue Gene machine based at the University of Groningen … to analyse data from 15,000 antennas spanning an area 350 km in diameter.

Protein Folding

Researchers at the University of Edinburgh, UK, are using a single-rack Blue Gene machine, capable of 6 trillion operations per second, to simulate protein folding and fluid mixing.

Climate Research

A single-rack Blue Gene is being shared by the National Center for Atmospheric Research and the University of Colorado, which runs simulations of ocean, weather and climate behaviour.

Cosmology

A team at the San Diego Supercomputing Center is using a single-rack Blue Gene to run Enzo — the centre's software that simulates how galaxies evolved from the Big Bang. The simulation, which has already predicted how the first stars formed, produced around 30,000 gigabytes of data on a typical recent run.

Drug Development

Japan's National Institute of Advanced Industrial Science and Technology is using its four-rack device to boost drug development. Researchers there hope to shed light on how drugs interact with their targets in the body.

Source: “Virtual Big Bangs and Digital Mushroom Clouds,” Nature Magazine, July 2005

“Killer Apps” on Blue Gene

Discipline	Application
Condensed Matter	CPMD, ddCMD, LSMS, ParaDis
Molecular Mechanics (ab initio)	Amber, CPMD, Qbox, abinit, CASTEP
Classical Molecular Dynamics	ddcMD, NAMD, GRASP, LAMMPS, MDCASK, NAMD, SPaSM, DL_POLY
BioInformatics/Life Science	Smith-Waterman, BLAST, mpiBLAST, pNeo
Neurosciences	Blue Brain (EPFL), SPLIT (KTH)
Astrophysics	FLASH, Enzo, Capreole
High Energy Physics	QCD
Weather / Climate	WRF, HOMME, POP
Plasma Physics	GTC, NIMROD
Computational Fluid Dynamics	AVBP, Miranda, NEK5000, Overflow, FUN3D, Raptor, LUDWIG
Nuclear Physics	QMC

Blue Gene Consortium

- **Purpose**

  A community of Blue Gene users who share experiences, help understand and develop the application space, assist in application porting and tuning and provide feedback on functional requirements for next generation

- **Membership Criteria**

  Membership is open to may be from any geography and any industry or organization, committed to provide a skill or undertaking such as performance measurements, application porting, OS/compiler testing

- **Consortium Facilitator**

  Argonne National Lab (ANL)
  http://www-fp.mcs.anl.gov/bgconsortium/default.htm

BG/L Consortium Members

LABORATORIES:

- Ames National Lab/Iowa State University
- Argonne National Laboratory
- Brookhaven National Laboratory
- Fermi National Laboratory
- Jefferson Laboratory
- Lawrence Berkeley National Laboratory
- Lawrence Livermore National Lab
- Oak Ridge National Laboratory
- Pacific Northwest National Laboratory
- Princeton Plasma Physics Laboratory

UNIVERSITIES:

- Boston University
- California Institute of Technology
- Columbia University
- Cornell University
- DePaul University
- Harvard University
- Illinois Institute of Technology
- Indiana University
- Iowa State University
- Louisiana State University
- Massachusetts Institute of Technology

UNIVERSITIES (cont'd):

- National Cntr for Atmospheric Research
- New York University - Courant Institute
- Northern Illinois University
- Northwestern University
- Ohio State University
- Pennsylvania State University
- Pittsburgh Super Computing Center
- Princeton University
- Purdue University
- Rutgers University
- Stony Brook University
- Texas A&M University
- University of California – Irvine
- University of California - San Francisco
- University of California - SDSC
- University of Chicago
- University of Colorado - JILA
- University of Delaware
- University of Hawaii
- University of Illinois Urbana-Champaign
- University of Minnesota
- University of North Carolina

UNIVERSITIES (cont'd):

- University of S California - ISI
- University of Texas at Austin - TACC
- University of Utah
- University of Wisconsin

INDUSTRY:

- Engineered Intelligence Corporation
- Gene Network Sciences
- IBM

INTERNATIONAL:

- Allied Engineering Corporation
- AIST
- ASTRON/LOFAR
- CERT
- EPFL
- John von Neumann Inst. for Computing
- National University of Ireland/ICHEC
- NIWS Co., Ltd.
- Trinity College, Trinity Centre for HPCC
- University of Edinburgh, EPCC
- University of Tokyo

Blue Gene Collateral

Specifications Sheet

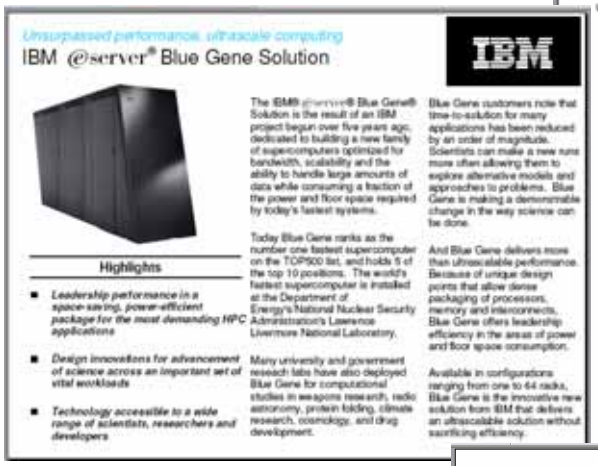
Unsurpassed performance, ultrascale computing

IBM @server® Blue Gene Solution

Highlights

- *Leadership performance in a space-saving, power-efficient package for the most demanding HPC applications*
- *Design innovations for advancement of science across an important set of vital workloads*
- *Technology accessible to a wide range of scientists, researchers and developers*

Brochure


IBM @server Blue Gene Solution

Redbooks

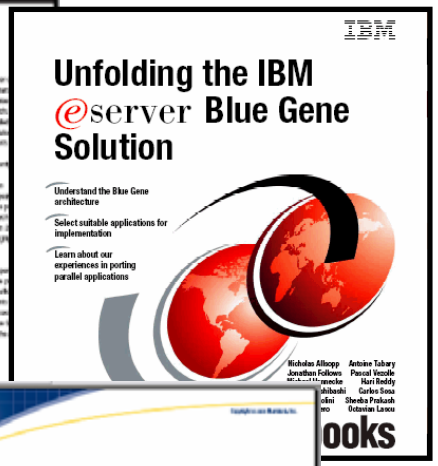
Unfolding the IBM @server Blue Gene Solution

Understand the Blue Gene architecture

Select suitable applications for implementation

Learn about our experiences in porting parallel applications

Application Notes Research Papers

- *Porting Applications to the IBM eServer Blue Gene/L System Solution*
- *Massive Parallel Quantum Computer Simulator*
- *Early Experience with Scientific Applications on the Blue Gene/L Supercomputer*
- *Scaling Physics and Material Science Applications on a Massively Parallel Blue Gene/L System*

TCO Study

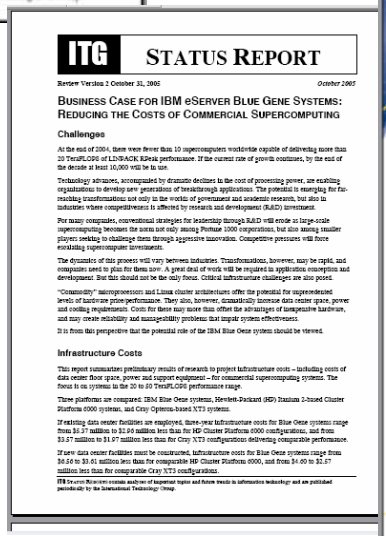
ITG STATUS REPORT

Business Case for IBM eServer Blue Gene Systems: Reducing the Costs of Commercial Supercomputing

Challenges

Infrastructure Costs

Analyst Report


Blue Gene's Teraflop Attack

Quick Note

What applications are currently running on Blue Gene and what is their performance?

Applications Performance on Blue Gene

- **BG is first HPC system to break barrier of 100+ TeraFlop/s sustained performance on real applications (Molecular Dynamics)**
  - **ddcMD – 101.5 TeraFlop/s (*7 hrs of Uranium atoms on 64 racks)***
  - **CPMD – 110.4 TeraFlop/s**
- **Several other applications have achieved two orders of magnitude or more higher performance than previously possible – successful scaling achieved from 1K to 100K processors**
- **Gordon Bell Prize competition at SC 2005**
  - **4 of 6 finalists based on Blue Gene**
  - **LLNL/IBM team won for "100+ TFlop Solidification Simulations on Blue Gene/L"**

ddcMD

Classical MD

2005 Gordon Bell Prize Winner

- **Scalable, general purpose code for performing classical molecular dynamics (MD) simulations using highly accurate MGPT potentials**
- **MGPT semi-empirical potentials, based on a rigorous expansion of many body terms in the total energy, are needed in to quantitatively investigate dynamic behavior of transitions metals and actinides**
- **Visualization of important scientific findings already achieved on BG/L: Molten Ta at 5000K demonstrates solidification during isothermal compression to 250 GPa**

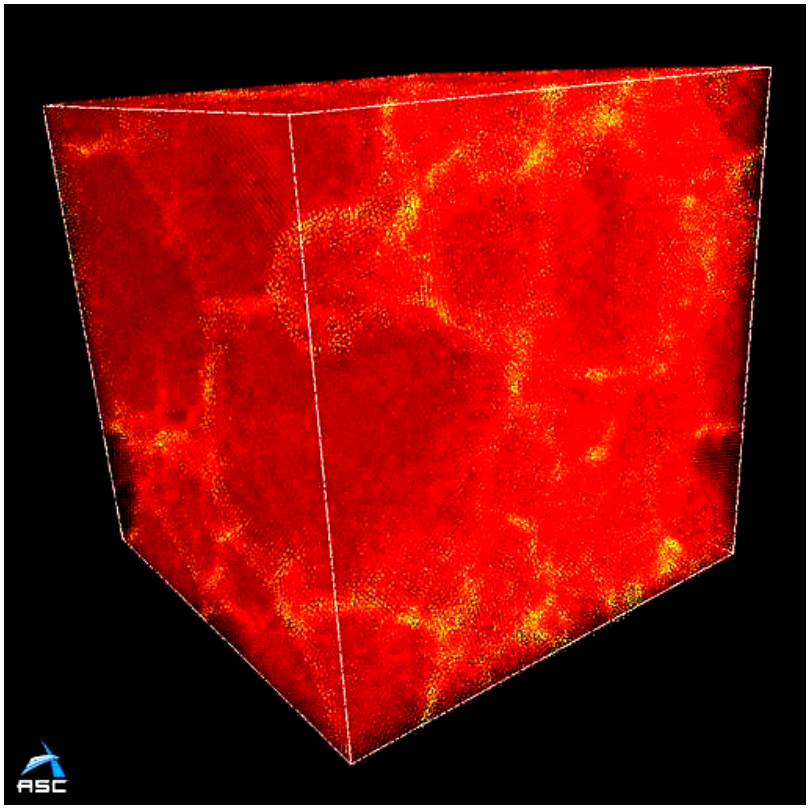

524 million atom simulations on 64K nodes are orders of magnitude larger than any previously attempted runs; *superb strong and weak scaling* expected for full machine - ("very impressive machine" says PI)

Performance of ddcMD on Blue Gene

Weak scaling: MGPT **Uranium** and Tantalum

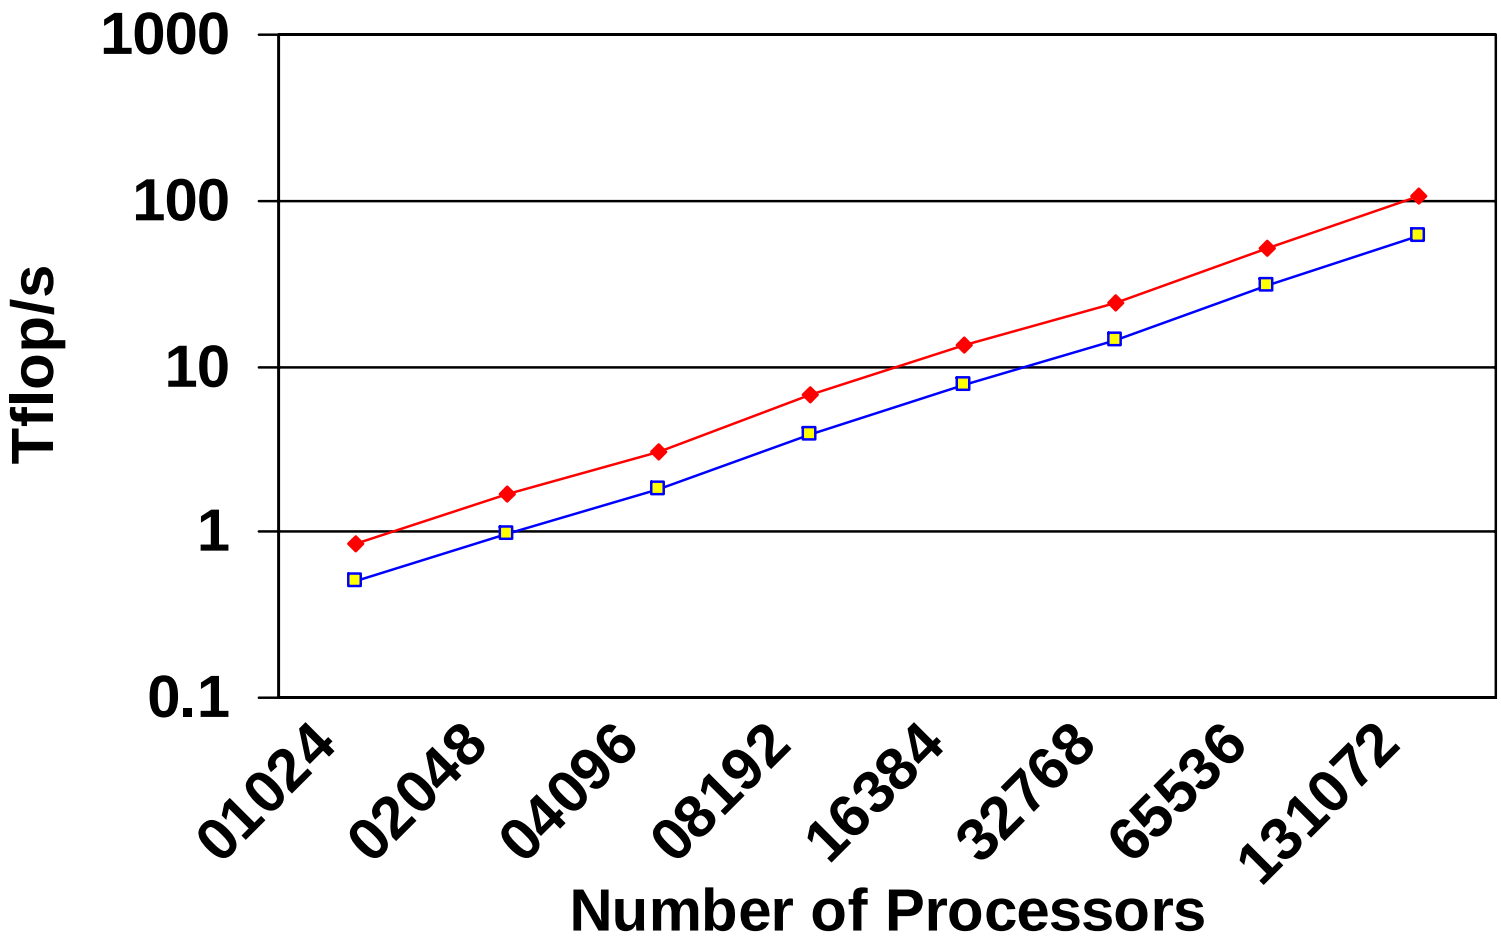

CPMD

Alessandro Curioni, Salomon Billeter, Wanda Andreoni

CPMD Performance on BG/L

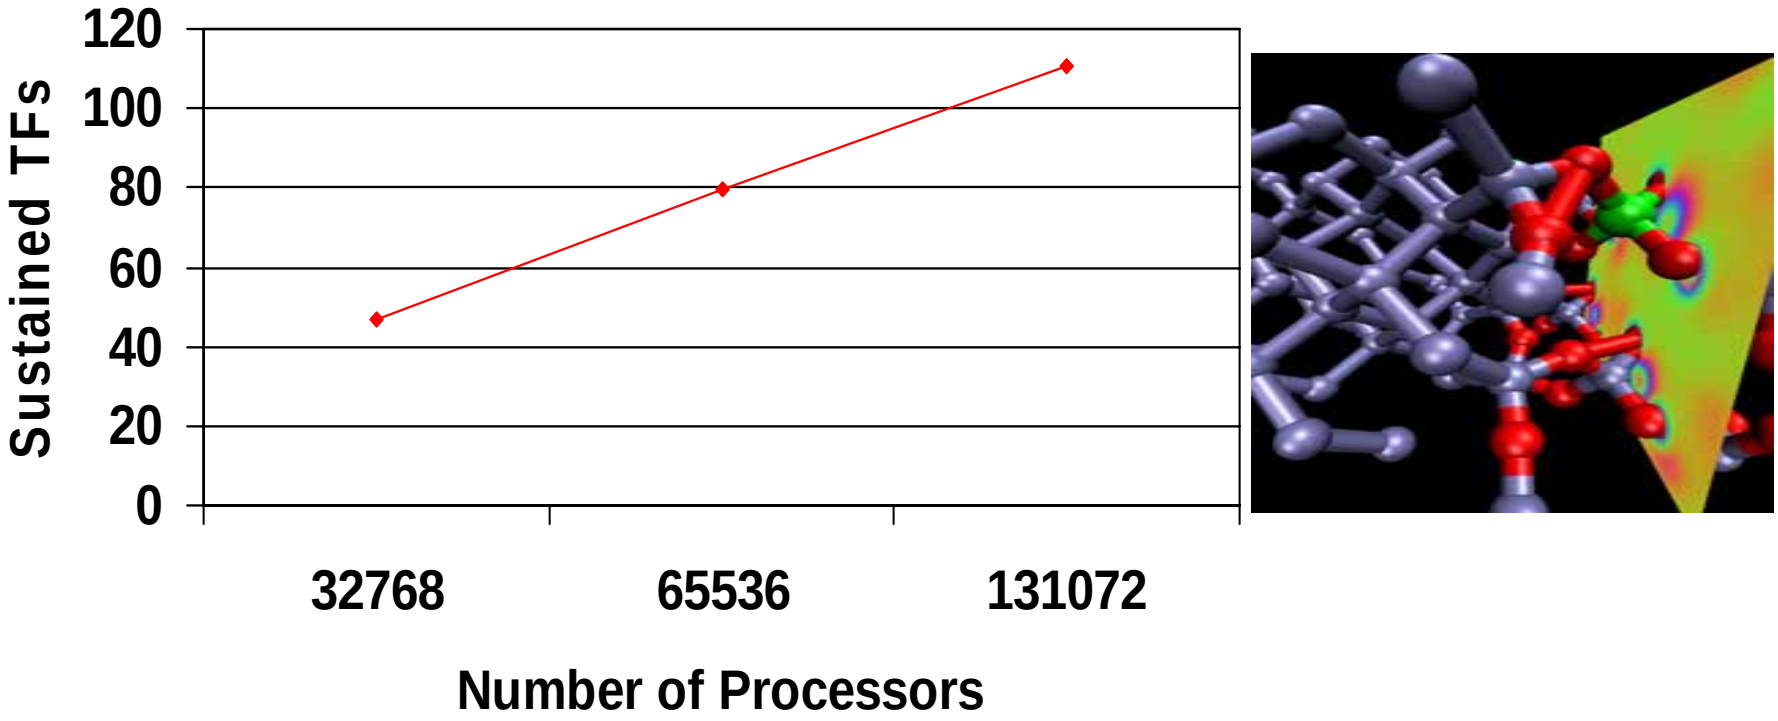

Developed at IBM Zurich from Car Parinello code
Uses Plane Wave Basis functions, FFT, MPI_Collectives
Ongoing project : IBM/LLNL Pd:H (~900 atoms) Hydrogen Storage
Achieved 110.4 Teraflop/s sustained on 64 racks BG/L (excellent strong scaling)

Conclusions

- Blue Gene represents an innovative way to scale to multi-TF/s
  - ✓ Leadership performance and price/performance
  - ✓ Massive scalability
  - ✓ Efficient packaging drives low power, cooling and floor space requirements
- Blue Gene balances massive scale-out capacity and preservation of familiar programming environments
- Blue Gene is applicable to a variety of computationally intensive workloads
  - ✓ High-energy physics
  - ✓ Astrophysics
  - ✓ Molecular dynamics and modeling
  - ✓ Climate modeling
  - ✓ Risk management and portfolio optimization
- Blue Gene technology is easily accessible for porting and optimization, including availability in IBM's On Demand Center
- IBM's ongoing R&D commitment ensures the viability and vitality of the Blue Gene solution

Read more about it…

http://www.ibm.com/servers/deepcomputing/bluegene.html

http://www-fp.mcs.anl.gov/bgconsortium/default.htm

http://www.research.ibm.com/journal/rd49-23.html

http://www.top500.org