




# Mit szól hozzá Watson - mitől szuper egy számítógép?

**Kohán Márk**

Systems Sales Representative





# IBM innovációs mérföldkövek




1944: Mark I



1948: SSEC



1956: RAMAC



1957: FORTRAN



1966:  
One-Device  
Memory Cell



1967:  
Fractals



1970: Relational  
Database



1971: Speech  
Recognition



1973:  
Winchester Disk



1979: Thin Film  
Recording Heads



1980:  
RISC



Nobel Prizes




1994:  
SiGe



1993: RS/6000 SP  
1996,97: Deep Blue




1997: Copper  
Interconnect Wiring




1998:  
Silicon-on-Insulator




1998:  
Microdrive



2001:  
Nanotube Transistor



2002: Millipede




2002:  
Molecule Cascade  
Logic Circuit




2004: Blue Gene/L  
The fastest supercomputer  
in the world

# 2008-ban regisztrált USA szabadalmak eloszlása


<http://www.ificlaims.com/IFIPatents010909.htm>






**IBM 4.186**




# A szuperszámítógépek tipikus alkalmazási területei



- Meteorológiai előrejelzések, szimulációk
- Blue Brain - Az emberi agy funkcionális működésének modellezése
- Watson – természetes nyelvek megértése
- Gyógyszerkutató, molekuláris biológia
- Folyadékok és gázok áramlásának szimulációja
- Adatbányászat




# Watson project - Jeopardy



Az emberi gondolkodásmóddal felruházott gép



# IBM Deep Thunder project





23-Oct-2005 - 20:00 EDT

Surface Total Precipitation and Winds  
Cloud Water Density at  $1.0e-04$  kg/kg a



- IBM p575 cluster - a számítási feladatokra
- IBM BladeCenter a vizualizáció kiszolgálására



# Turboinstitute - Kis méretű vízturbinák tervezése



# A feladvány: a hatékonyság növelése szimulációk segítségével



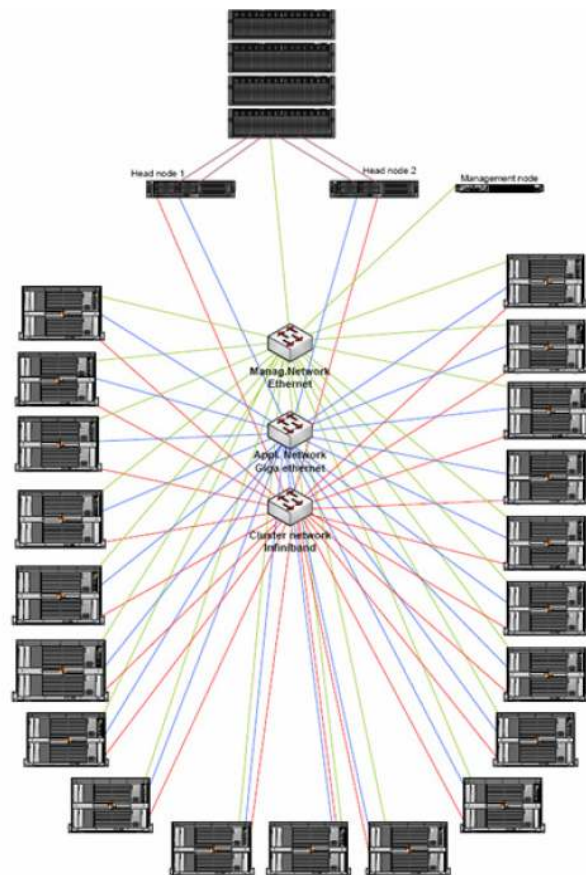

Computational Fluid Dynamics

Navier-Stokes

$$\rho \frac{\partial \vec{v}}{\partial t} = \rho \vec{F} - \nabla p - \rho \vec{v} \nabla \vec{v} + \mu \nabla^2 \vec{v}$$




# A HPC infrastruktúra - IBM BladeCenter és X3850 HPC cluster














- 105KW
- 3,5 tonna
- 4.1TB RAM
- 2048 CPU mag
- 10TB háttértár
- 36TFLOPS




# Knorr-Bremse kutatás-fejlesztés



1. Fizikai kölcsönhatások modellezése
2. Folyadék dinamikai analízis és optimalizáció
3. Hőáramlás analízis
4. Folyadék szerkezet és kölcsönhatások
5. Numerikus optimalizációk
6. Alkatrészek életciklusának analízálása




# Egy kis IBM szuperszámítógép történelem




1954	IBM NORC	67 kOPS	Department of Defense U.S. Naval Proving Ground, Dahlgren, Virginia
1958	IBM AN/FSQ-7	400 kOPS	25 U.S. Air Force sites across the continental USA and 1 site in Canada
1961	IBM 7030 "Stretch"	1.2 MFLOPS	AEC-Los Alamos National Laboratory, New Mexico, USA
1997	IBM DeepBlue	11.38 GFLOPS	The Deep Blue chess computer which defeated Kasparov
2000	IBM ASCI White	7.226 TFLOPS	DoE-Lawrence Livermore National Laboratory, California, USA
2004	IBM Blue Gene/L	70.72 TFLOPS	DoE/IBM Rochester, Minnesota, USA
2005	IBM Blue Gene/L	136.8 TFLOPS	National Nuclear Security Administration, Lawrence Livermore National
2006	IBM Blue Gene/P	280.6 TFLOPS	
2007	IBM Blue Gene/P	478.2 TFLOPS	
2008	IBM Cell	1.026 PFLOPS	IBM Roadrunner DoE-Los Alamos National Laboratory, New Mexico


# 1954 - IBM NORC - 67 kOPS



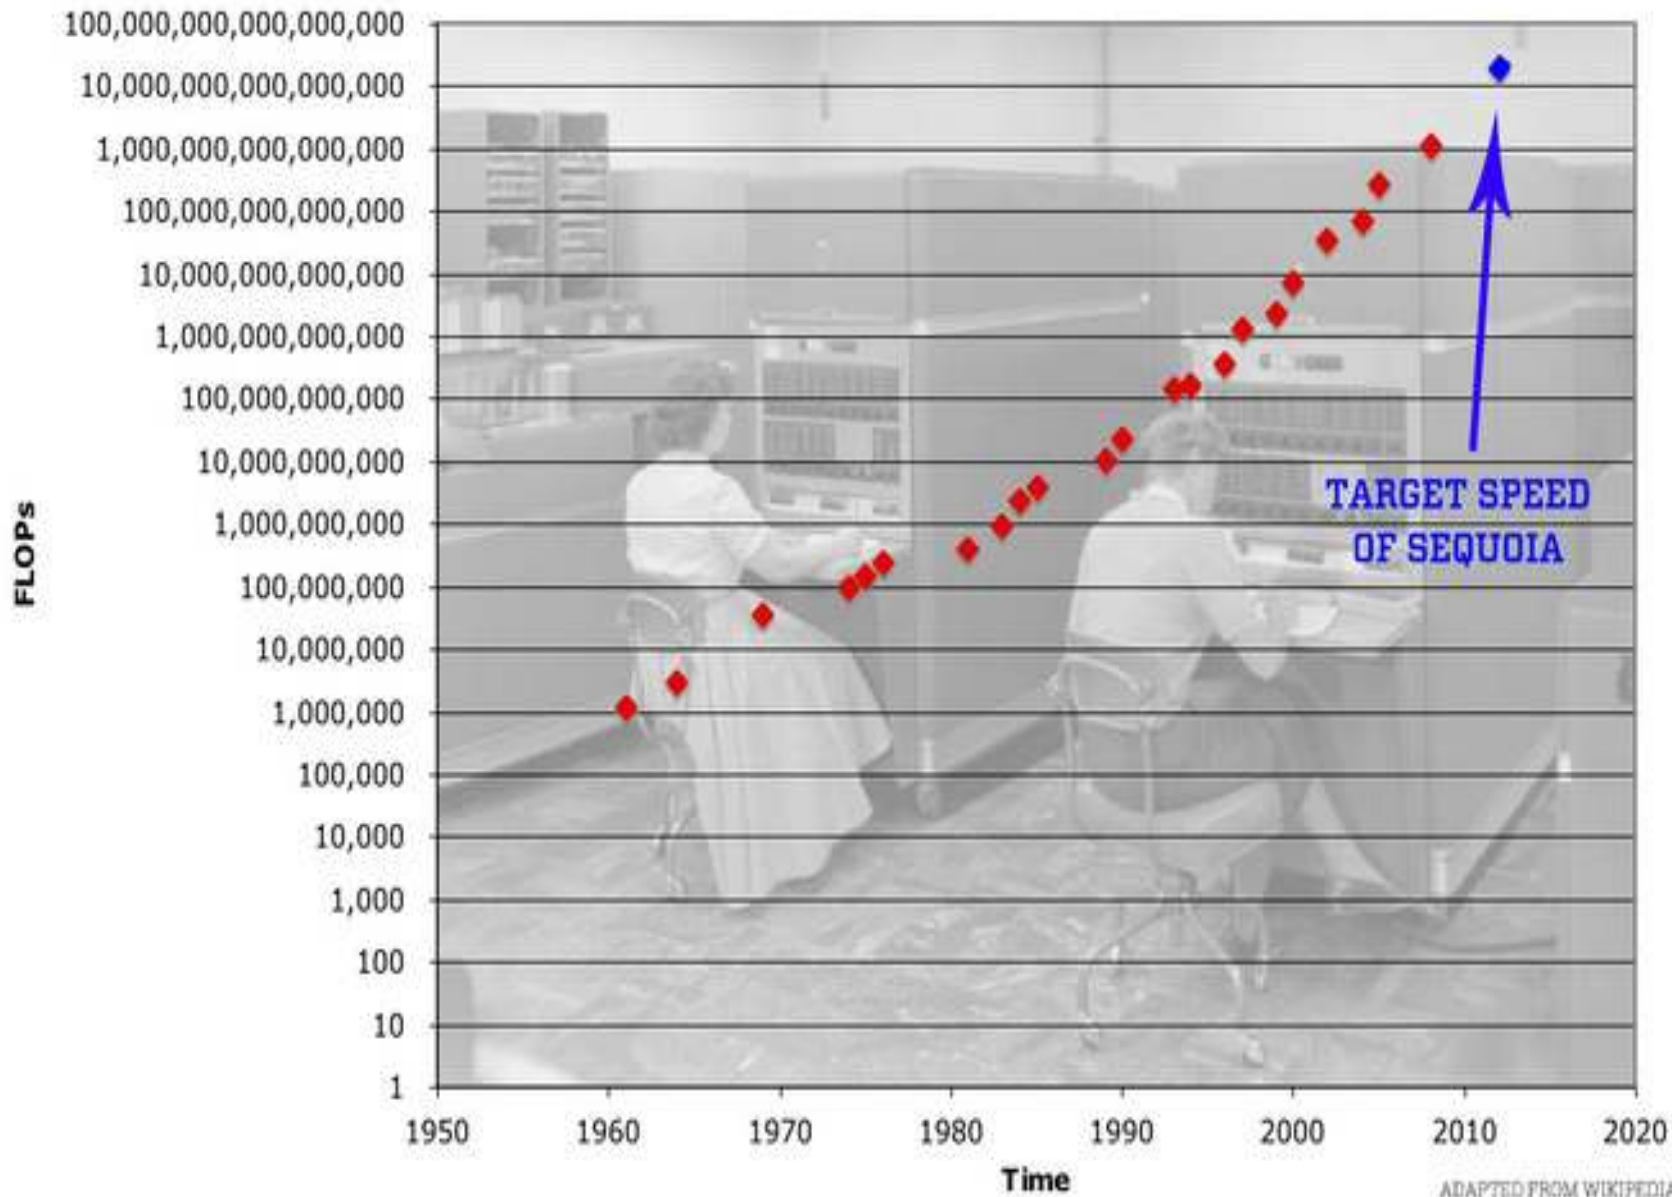

# 2008 - IBM Roadrunner - 1.026 PFLOPS



# 2012 - Sequoia Project – 20PFLOPS






# HISTORICAL PERFORMANCE OF THE WORLD'S FASTEST SUPERCOMPUTERS






# Growth in Supercomputer Power

Logarithmic Plot








## ...Mi is történik a TOP500 lista élén?



Rank	Site	Computer/Year Vendor	Cores	$R_{max}$	$R_{peak}$	Power
1	DOE/NNSA/LANL United States	Roadrunner - BladeCenter QS22/LS21 Cluster, PowerXCell 8i 3.2 Ghz / Opteron DC 1.8 GHz, Voltaire Infiniband / 2008 IBM	129600	1105.00	1456.70	2483.47
2	Oak Ridge National Laboratory United States	Jaguar - Cray XT5 QC 2.3 GHz / 2008 Cray Inc.	150152	1059.00	1381.40	6950.60
3	Forschungszentrum Juelich (FZJ) Germany	JUGENE - Blue Gene/P Solution / 2009 IBM	294912	825.50	1002.70	2268.00
4	NASA/Ames Research Center/NAS United States	Pleiades - SGI Altix ICE 8200EX, Xeon QC 3.0/2.66 GHz / 2008 SGI	51200	487.01	608.83	2090.00
5	DOE/NNSA/LLNL United States	BlueGene/L - eServer Blue Gene Solution / 2007 IBM	212992	478.20	596.38	2329.60

# A Green500.org oldalon található lista




Green500 Rank	MFLOPS/W	Site*	Computer*	Total Power (kW)	TOP500 Rank*
1	536.24	Interdisciplinary Centre for Mathematical and Computational Modeling, University of Warsaw	BladeCenter QS22 Cluster, PowerXCell 8i 4.0 Ghz, Infiniband	34.63	422
2	458.33	DOE/NNSA/LANL	BladeCenter QS22/LS21 Cluster, PowerXCell 8i 3.2 Ghz / Opteron DC 1.8 GHz, Infiniband	138.00	61
2	458.33	IBM Poughkeepsie Benchmarking Center	BladeCenter QS22/LS21 Cluster, PowerXCell 8i 3.2 Ghz / Opteron DC 1.8 GHz, Infiniband	138.00	62
4	444.94	DOE/NNSA/LANL	BladeCenter QS22/LS21 Cluster, PowerXCell 8i 3.2 Ghz / Opteron DC 1.8 GHz, Voltaire Infiniband	2483.47	1
5	428.91	National Astronomical Observatory of Japan	GRAPE-DR accelerator Cluster, Infiniband	51.20	277
6	371.67	ASTRON/University Groningen	Blue Gene/P Solution	94.50	124
7	371.67	IBM - Rochester	Blue Gene/P Solution	126.00	84
7	371.67	IBM Thomas J. Watson Research Center	Blue Gene/P Solution	126.00	85
7	371.67	Max-Planck-Gesellschaft MPI/IPP	Blue Gene/P Solution	126.00	86
7	371.67	Bulgarian State Agency for Information Technology and Communications (SAITC)	Blue Gene/P Solution	63.00	245
7	371.67	Moscow State University	Blue Gene/P Solution	63.00	246
7	371.67	Oak Ridge National Laboratory	Blue Gene/P Solution	63.00	247
7	371.67	Stony Brook/BNL, New York Center for Computational Sciences	Blue Gene/P Solution	63.00	248
14	368.89	EDF R&D	Blue Gene/P Solution	252.00	36
15	368.30	IDRIS	Blue Gene/P Solution	315.00	24
16	367.40	King Abdullah University of Science and Technology	Blue Gene/P Solution	504.00	14
17	366.58	DOE/NNSA/LLNL	Blue Gene/P Solution	1134.00	9
18	363.98	Forschungszentrum Juelich (FZJ)	Blue Gene/P Solution	2268.00	3
19	363.98	Argonne National Laboratory	Blue Gene/P Solution	1260.00	7
20	273.06	HWW/Universitaet Stuttgart	NEC HPC 140Rb-1 Cluster, Xeon X5560 2.8Ghz, Infiniband	186.00	77
21	271.52	NASA/Goddard Space Flight Center	iDataPlex, Xeon X55xx QC 2.8 GHz, Infiniband	143.60	101

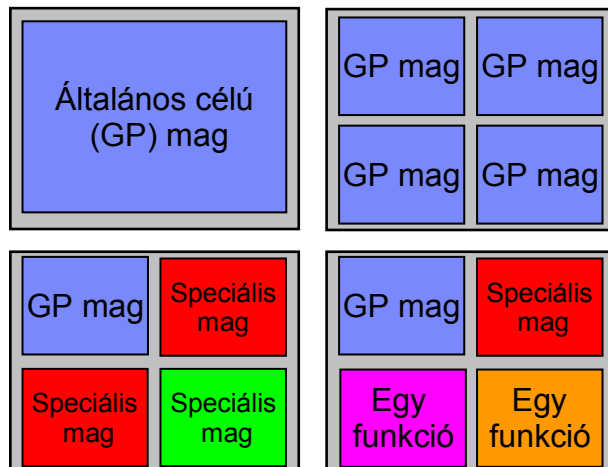
# A teljesítmény növelésének nehézségei



## Készen állunk-e a masszív párhuzamosításra?



- Hogyan tudjuk a jelenlegi és jövőbeni technológiát a legjobban kihasználni?




**Párhuzamos gondolkodás  
és problémamegoldás**

**Párhuzamos programozási modellek  
és eszközök**

**Párhuzamos számítógépek**

- Nagyobb akadályokkal kell megküzdenünk, mint a puszta technológia...

# A teljes kép




Szoftver

Eszközök

Programozási modellek

Alkalmazások

x86



HPCs



BLUE GENIE



MAXIMIZE THE POWER OF THE CELL/B.E. PROCESSOR



Infrastruktúra


Hálózat





Tárolás




GPFS




# A Blue Gene /P felépítése





**Blue  
Brain  
Project**




**Henry Markram**




**Felix Schürmann**




**Sean Hill**



## PowerXCell 8i architektúra




- Hibrid, 9 magos mikroprocesszor
  - PPE (vezérlés) + 8 SPEs (számításra optimalizált)
- 64 bites architektúra
  - OS, SW fejlesztés, hagyományos alkalmazások
- Nagy sávszélesség
  - 155+ konkurens memória-processzor tranzakció
- Energiahatékonyság
  - 460 GFLOPS (SP) / 217 GFLOPS (DP) - 92 Watt
- Real-time
  - Erőforrás-hozzárendelés
- Biztonsági támogatás
  - Minden SPE biztonságos processzorként konfigurálható







## Az IBM BladeCenter QS22 paramétereit




<b>Processzor</b>	2x 3.2GHz Cell/B.E processzor - 1x PPE (Power Proc. Element)/proc - 8x SPE (Sysnergic Proc.Element)/proc
<b>L2 cache</b>	512KB/proc + 256KB/SPE
<b>Memória</b>	32GB (16 GB/proc)
<b>I/O buffer</b>	4x 512MB VLP DIMM
<b>Belső diszk</b>	8 GB Flash drive (opció)
<b>LAN</b>	Dual GbE interfész
<b>Nagy sebességű kapcsolat</b>	Dual 4x Infiniband Interfész
<b>Felépítés</b>	Szabványos Blade (keskeny), max 14 QS22/Blade keret



# 3D Computer Tomografía @ MAYO



recon-anim-pc-cvid.avi



recon-anim-cell-cvid.avi

## Mayo Clinic and IBM Score Significant Advance in Real-Time Medical Imaging

*Imaging Application Produces Results up to Fifty Times Faster Than on Typical Processors*


ROCHESTER, MN and ARMONK, NY - 09 Apr 2007:

Collaborators from Mayo Clinic and IBM (NYSE: IBM) have exploited parallel computer architecture and memory bandwidth to dramatically speed up the processing of 3-D medical images. The advance significantly aids image registration -- the computer-enhanced alignment of two medical images obtained at different dates or by using different imaging devices, in three-dimensional space. With the images properly aligned over one another, a radiologist can more easily detect structural changes such as the growth or shrinkage of tumors.

The results will be presented in full in a joint presentation by Mayo Clinic and IBM at the IEEE (Institute of Electrical and Electronics Engineers) International Symposium on Biomedical Imaging in Washington, D.C., April 12-15.

<http://www-03.ibm.com/press/us/en/pressrelease/21342.wss>

# Az IBM Blade szerverek



**HS22**  
High-performance  
enterprise  
intel

**JS23**  
High-performance  
w/native  
virtualization  
IBM

**LS22**  
High-performance  
AMD

**LS42**  
Scalable, enterprise  
performance  
AMD


**QS22**  
High-performance  
IBM

**JS43**  
High-performance  
w/native  
virtualization  
IBM


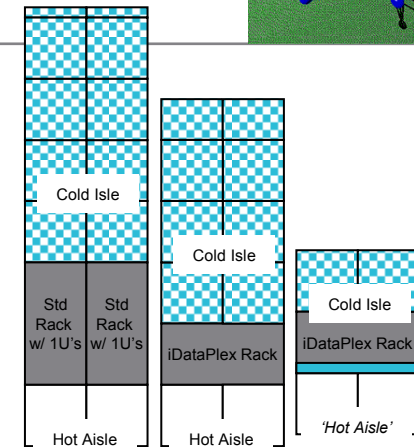

## Csak az IBM-nél:

- ✓ Közös, kompatibilis blade szerverek
- ✓ Közös, kompatibilis ipari szabvány I/O kapcsoló
- ✓ Közös menedzsment infrastruktúra
- ✓ Nincs „single point of failure” – minden redundáns
- ✓ Nyílt szabvány


# IBM System x iDataPlex



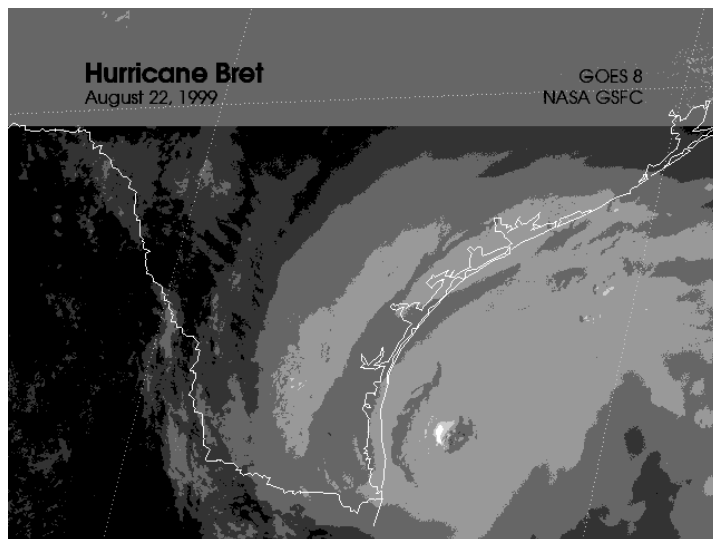
- Nagy elemsűrűség
- Integrált infrastruktúra
- HW redundancia minimális
- Vízhűtéses hátsó panel
- Újszerű rack elrendezés
- Előregyártott konfiguráció
- 29 kW, 6 TFLOPS




# NASA föld megfigyelő rendszer




- Atmoszféra kutatás
- Klíma szimulációk
- Hurrikánok kialakulásának és szerkezetének vizsgálata




# IBM Power Systems 575 HydroCluster – POWER6




*THE NEW POWER EQUATION*




- **Node méret**
  - 2U, 24" X 51"
  - 32 X 4.7GHz POWER6 mag
  - 256 GB memória
- **Rack adatok**
  - 42U, 24"
  - 14 node
  - 448 mag
  - 3.5 TB memory
  - 8 TFLOPS
- **Környezet**
  - AIX & Linux
  - Vízhűtés
  - Zöld adatközpont





✓ **Vízhűtés - energiahatékony**




Kaleidoscope project, kőolajkutatás,  
szeizmikus képalkotás, Cell/B.E.  
Repsol YPF, 3DGeo,  
Barcelona Supercomputing Center




Lipidmembrán szimulációja, molekuláris  
biológia, Blue Gene,  
Argonne Leadership Computing Facility



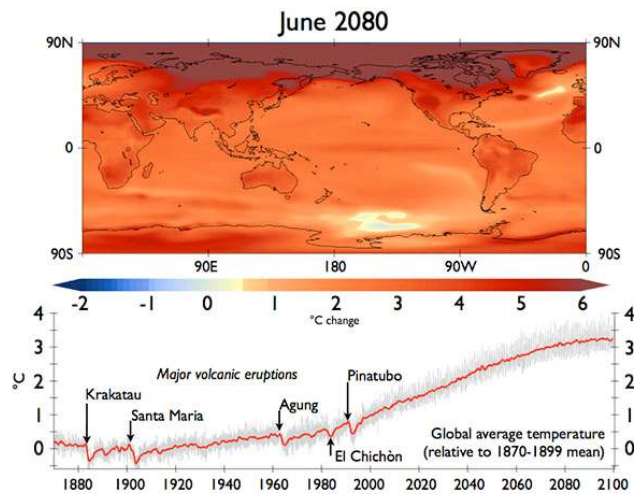

Molekuláris dinamika, fémek atomi szerkezetének  
alakulása, Blue Gene, Scalable Parallel Short-range  
Molecular Dynamics,  
Los Alamos National Laboratory



G-Protein Coupled Receptor  
 Gyógyszerkutatás, molekuláris biológia,  
 IBM Thomas Watson kutatóközpont,  
 Blue Matter keretrendszer  
 Blue Gene




Sejthártya-fehérjék szimulációja  
 IBM Research, Blue Gene




Klimatikus modellezés, „Blue Fire” Power 575,  
 National Center for Atmospheric Research







Szélcsatorna-modell  
folyadékok áramlása, szilárdtest-mechanika,  
végelem-módszerek, Navier-Stokes egyenlet  
Dortmundi Egyetem és Digital Medics

**CELL/B.E.**




Blue Brain Project, emberi agy funkcionális  
modellezése, neokortikális oszlopok  
szimulációja, EPFL, **Blue Gene**



Nukleáris fegyverzet öregedésének  
elemzése, DoE, LANL


**IBM RoadRunner**

## Európa jelentősebb installációi



- JUGENE  
(Jülichi Egyetem, Blue Gene – 180 TF)
- Mare Nostrum  
(Barcelona, BSC, Linux, JS20 PowerPC Blades)
- Nautilus  
(RoadRunner architektúra, Varsói Egyetem)





---

**Köszönöm a figyelmüket!**

