Computing at CERN - III

Summer Student Lectures 2002

Jamie Shiers

http://cern.ch/jamie

Lecture III

- Computing at CERN Today
- Software at CERN Today
- **The future & LHC Computing**

Homework

Review of homework from lecture II

Exercise II

- What will the CERN Computing environment look like in 10 years?
- Hint: some of the key elements exist today, albeit possibly in a different flavour.

Lecture III

- Computing at CERN Today
- Software at CERN Today
- **The future & LHC Computing**

- *"T ... ot*
  wi...

Willi ... erm
“Cy...

- Un... 1990
- Wh... ve?
- W... ?

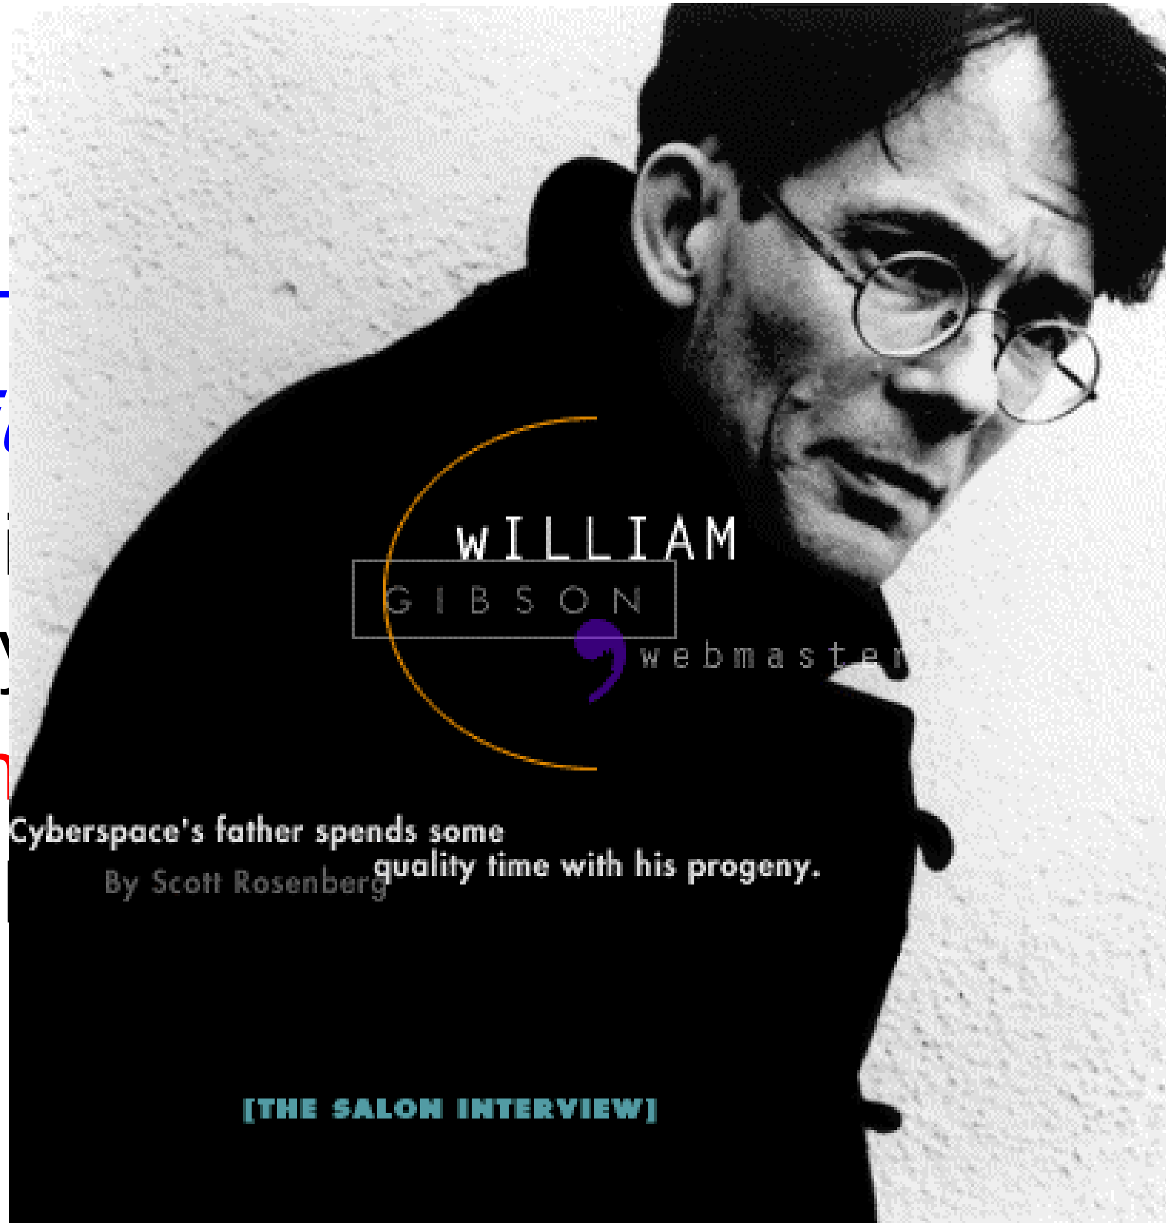

Predictions from 1945

- “As we may think”
  – Vannevar Bush
- Describes “memex”
  – A memex is a device in which an individual stores all his books, records, and communications, and which is mechanized so that it may be consulted with exceeding speed and flexibility. It is an enlarged intimate supplement to his memory.
- Used in much the same way as the Web

Lessons from the past

- ☠ Technologies explicitly designed to be the future rarely are...
  - Multics, ISO/OSI Network model, ADA, Alpha processor, Object Databases, Iridium, 3G, ...
- ☺ Very rapid advances in some areas
  - e.g. processor power, storage, ...
- ☹ Seemingly little in others...
  - Unix / Linux, Xerox PARC: Alto PC, Ethernet, distributed computing ... are all 1/4 century old!

Lessons from the past

- Technologies explicitly designed to be the future rarely are…
  - Multics, ISO/OSI Network model, ADA, Alpha processor, **Object Databases**, Iridium, 3G, …
- Very rapid advances in some areas
  - e.g. processor power, storage, …
- Seemingly little in others…
  - Unix / Linux, Xerox PARC: Alto PC, Ethernet, distributed computing … are all 1/4 century old!

- Re
  -
- Co 90s
  - .
- Go *ge and co data ve* (Cattell)
  - ring & Control,
  - al” DBMS
- Sta
- Pre **lace** RDBMS

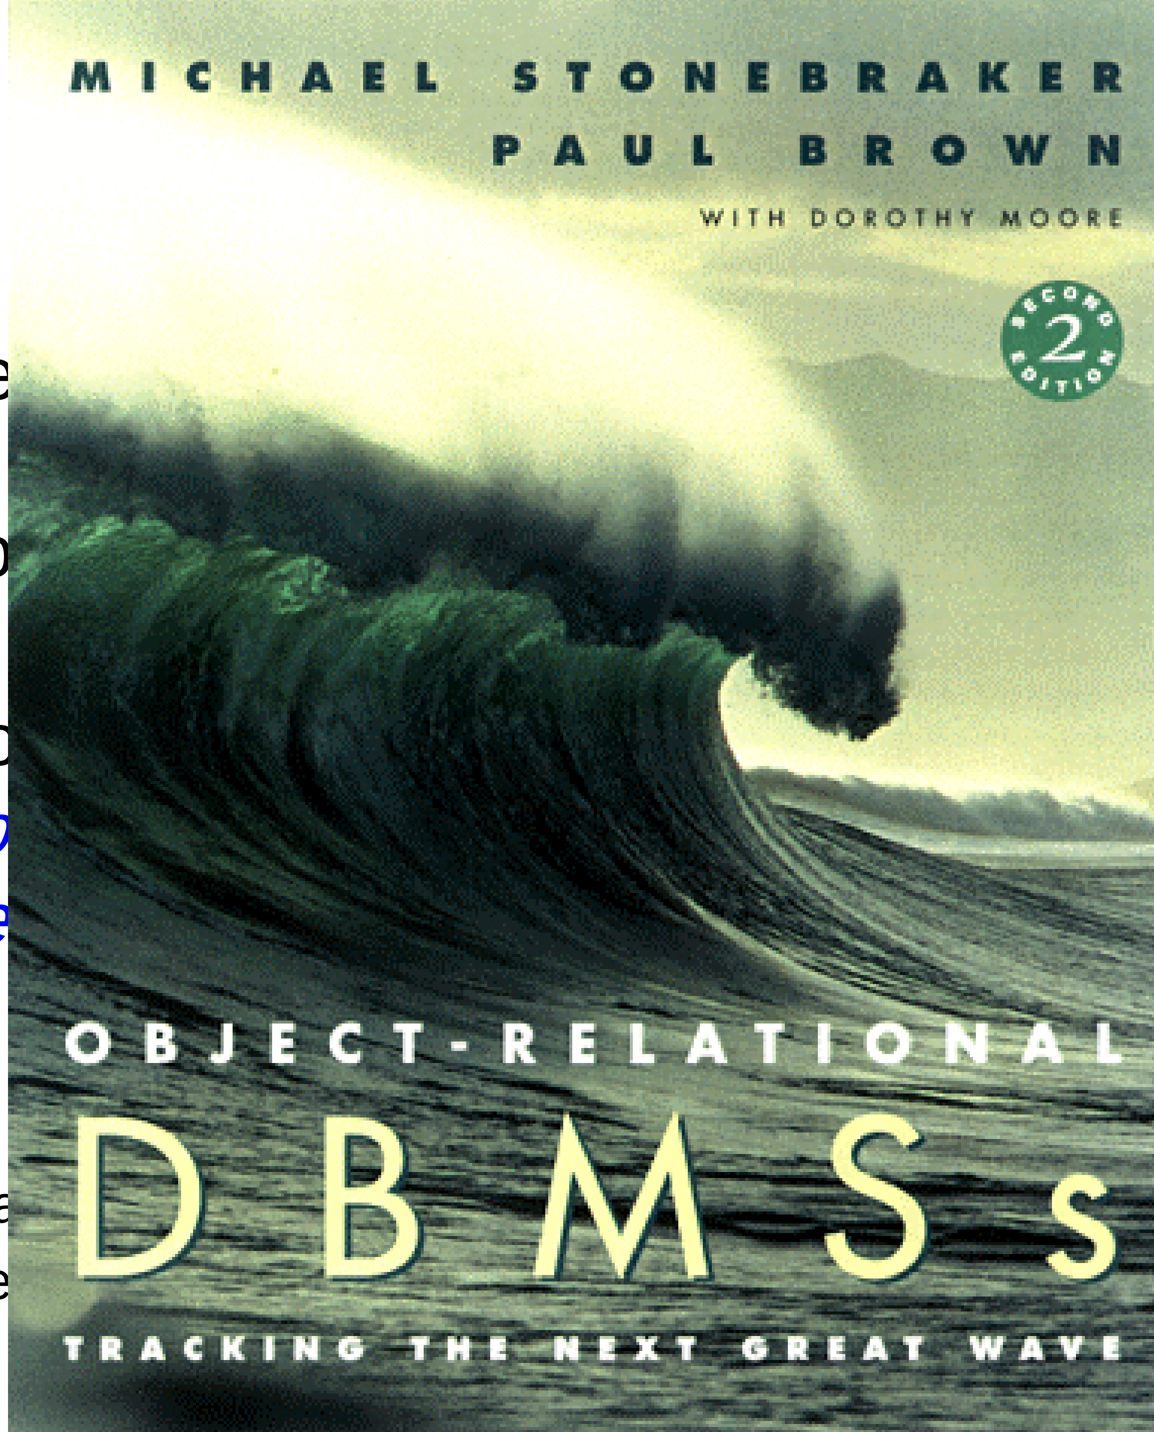

Lessons from the past

- ☠ Technologies explicitly designed to be the future rarely are…
  - Multics, ISO/OSI Network model, ADA, Alpha processor, Object Databases, **Iridium**, 3G, …
- ☺ Very rapid advances in some areas
  - e.g. processor power, storage, …
- ☹ Seemingly little in others…
  - Unix / Linux, Xerox PARC: Alto PC, Ethernet, distributed computing … are all 1/4 century old!

The Future...

- Planning for the future:
  - Necessarily conservative: basically extrapolations of current / immediate technology
- Predicting the future:
  - Much more speculative ... and fun...

The Future's Here..

- Key predictions of Telecom 1999:
  - ✓ Convergence of mobile phones & PDA
    - Phones with main PDA apps built-in exist
    - Phones with full PDA functionality too...
  - ✗ Emergence of 3G networks
    - Lack of clear "killer app"
    - Down-loading ring-tones is clearly not it
  - ➔ Wireless networks offer strong competition

April Fool's Day…

- More computing power than the Apollo space programme…

W ... s...

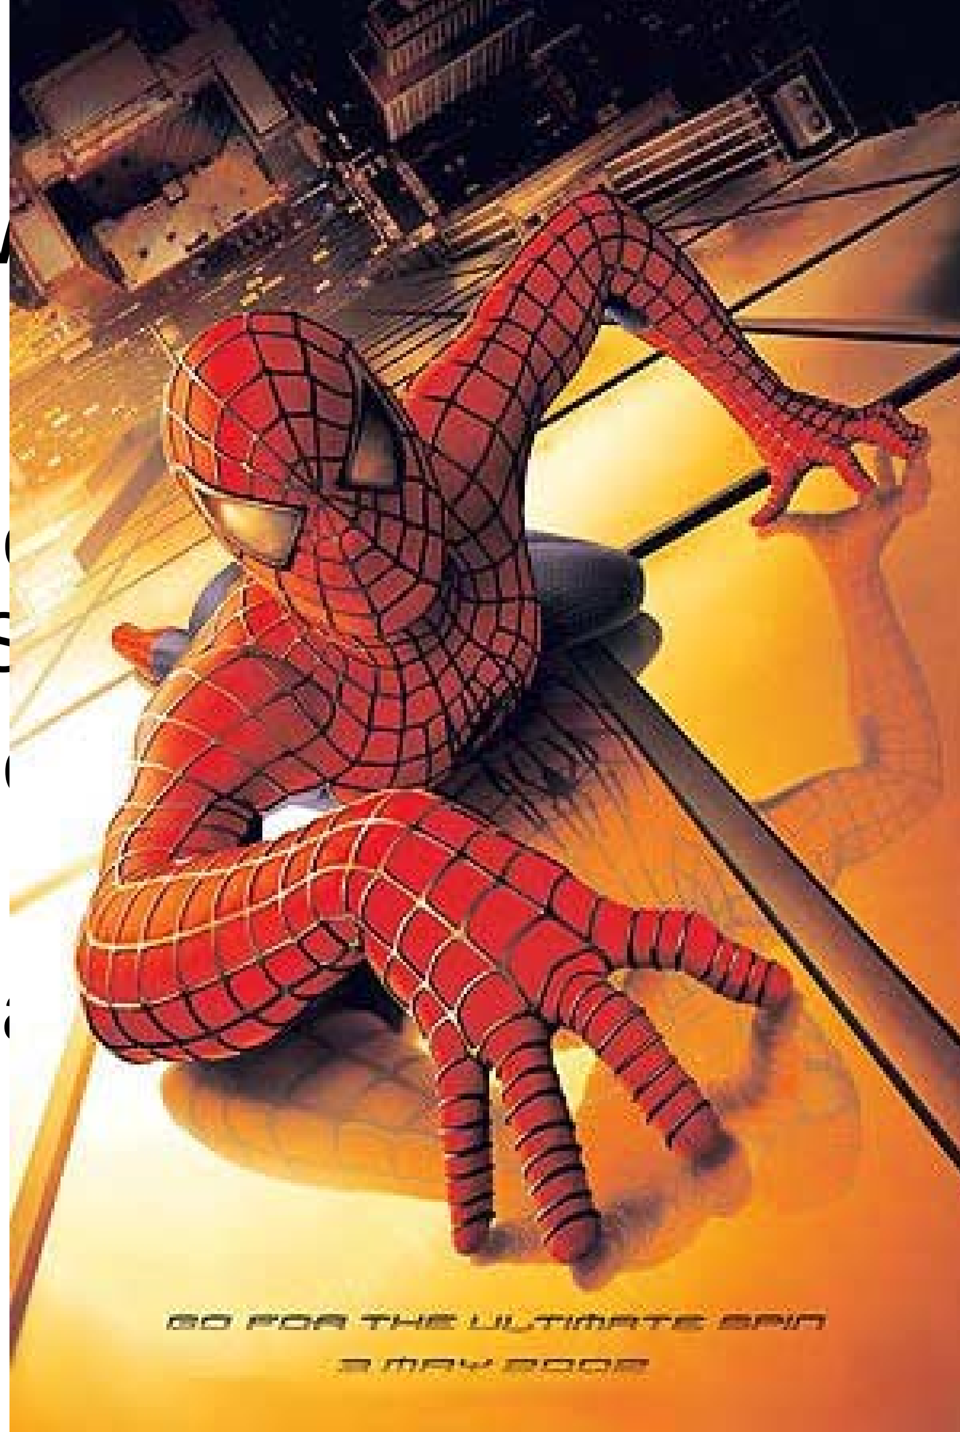

- No ... ns such as S
- No
- No
- No a ... tions...

LCG

Requirements per LHC Experiment

Processor power	> 10^6 SPECint95
Data volume	> 2PB / year
Data rate	> 1Tbit / second
# addressable objects	> 10^9
# users	10^3
# data traversals	10 - 10^2

Few GB/s per PB

HEP Computing Characteristics

- Large numbers of independent events
  - ✓ trivial parallelism
- Large data sets
  - ✓ smallish records; mostly read-only
- Modest I/O rates
  - ✓ few MB/sec per fast processor
- Modest floating point requirement
  - ✓ SPECint performance

☹ Very large **aggregate** requirements

Cost Estimates for CERN

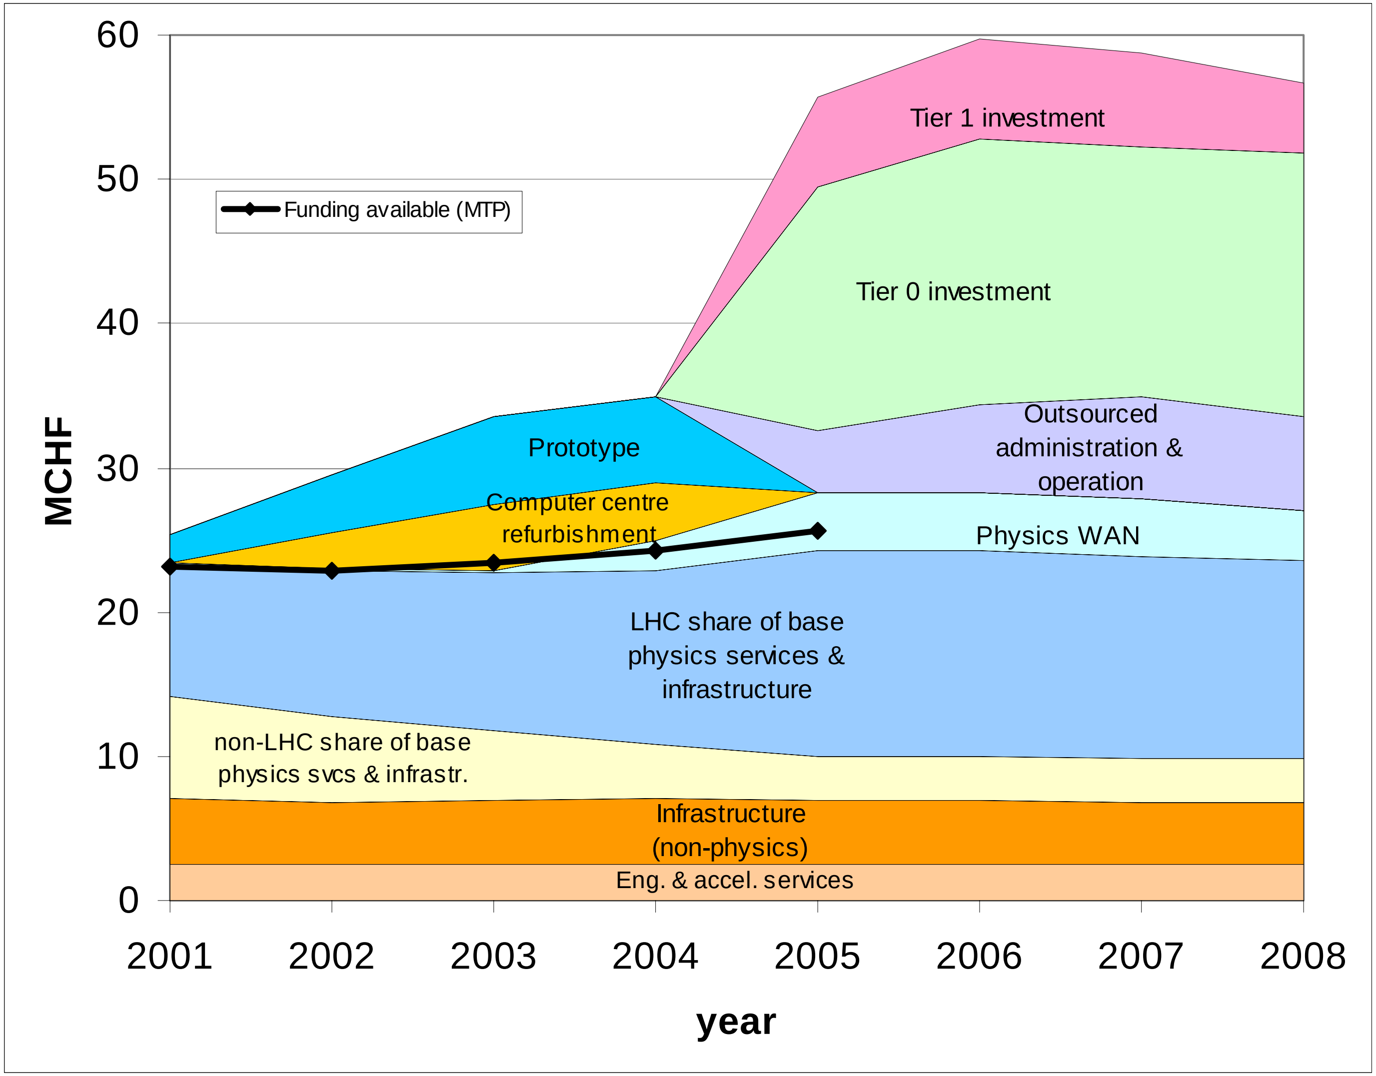

Evolution of LHC Prototype

Capacity							
		year	***2001***	***2002***	***2003***	***2004***	***2005***
processor farm							
no. of 2-cpu systems installed			182	400	400	600	800
estimated total capacity (SI95)			15000	33000	33000	69800	121800
disk storage							
no. of disks installed			200	480	480	960	1600
estimated total capacity (TB)			15	47	47	143	271
tape drives							
total capacity (achievable MB/sec)			150	350	450	600	800
automated media							
total capacity (TB)			30	100	200	400	600

PASTA

CERN Technology Tracking for the LHC

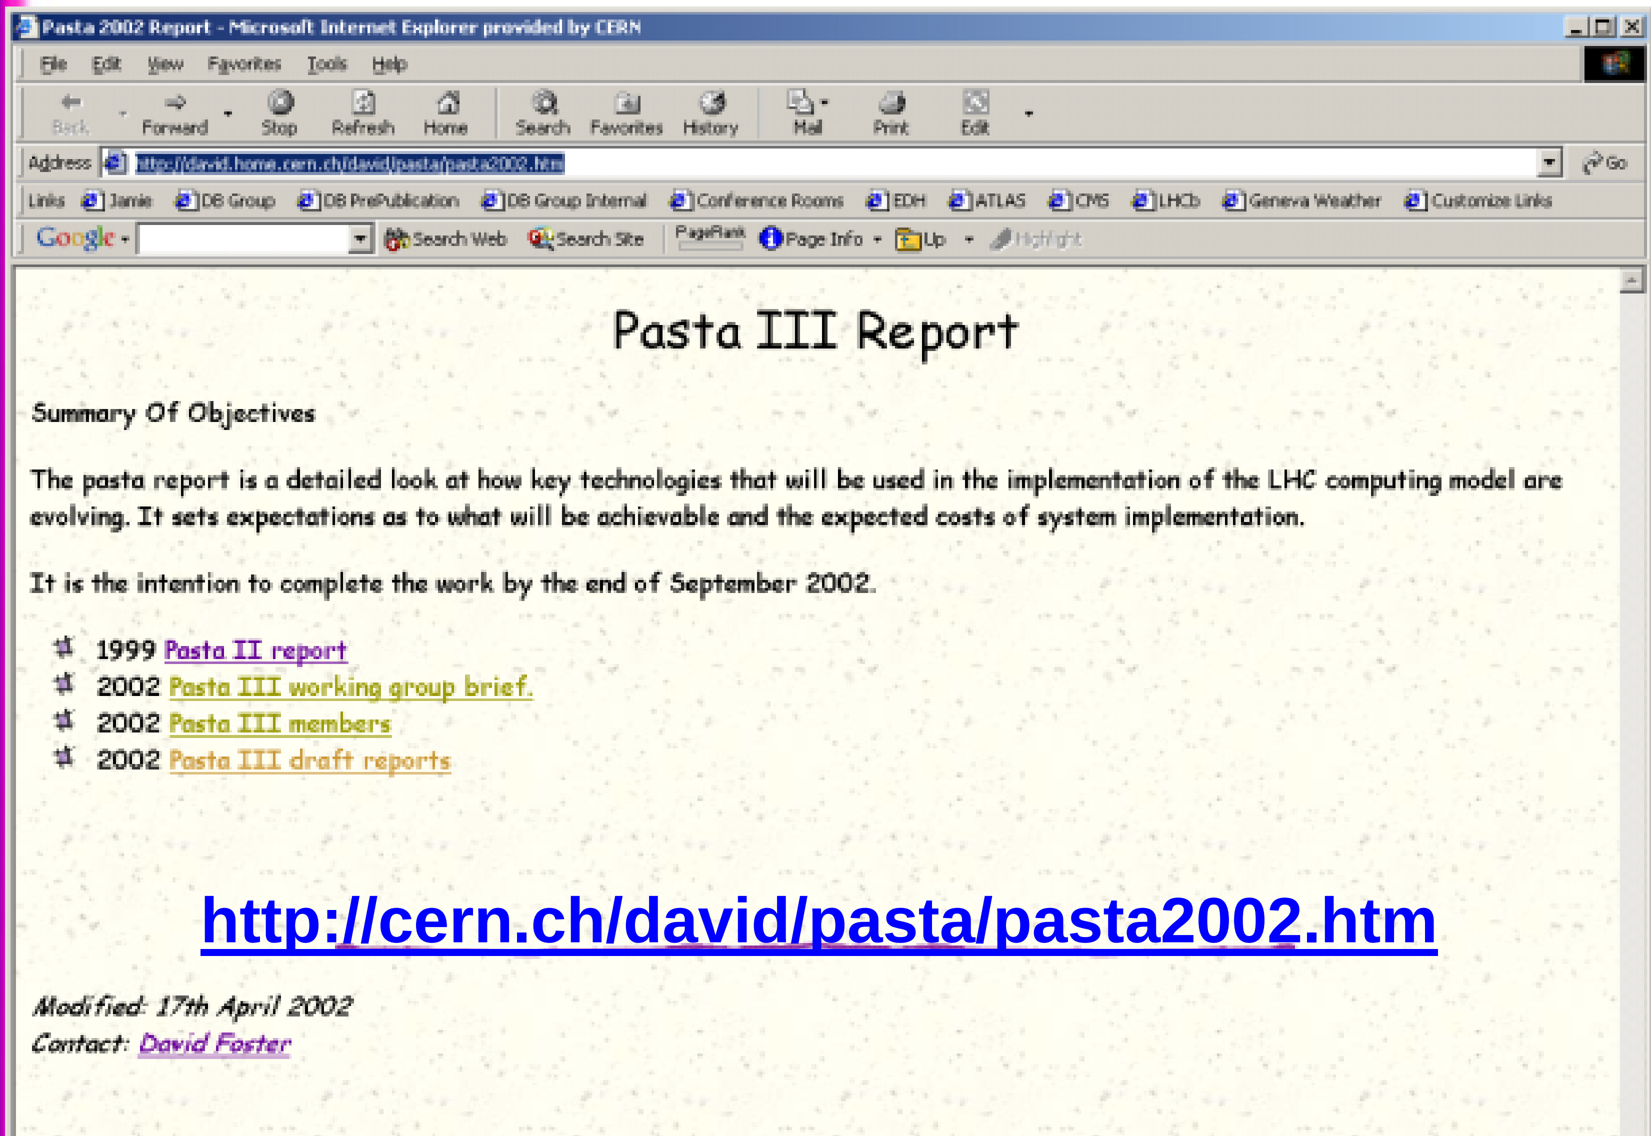

Storage Predictions

Storage Colloquium

- Wednesday 7th August, 14:00, main auditorium
- Jai Menon, IBM Storage Research
  - Storage Tank, IceCube

LHC – A Multi-PB Problem!

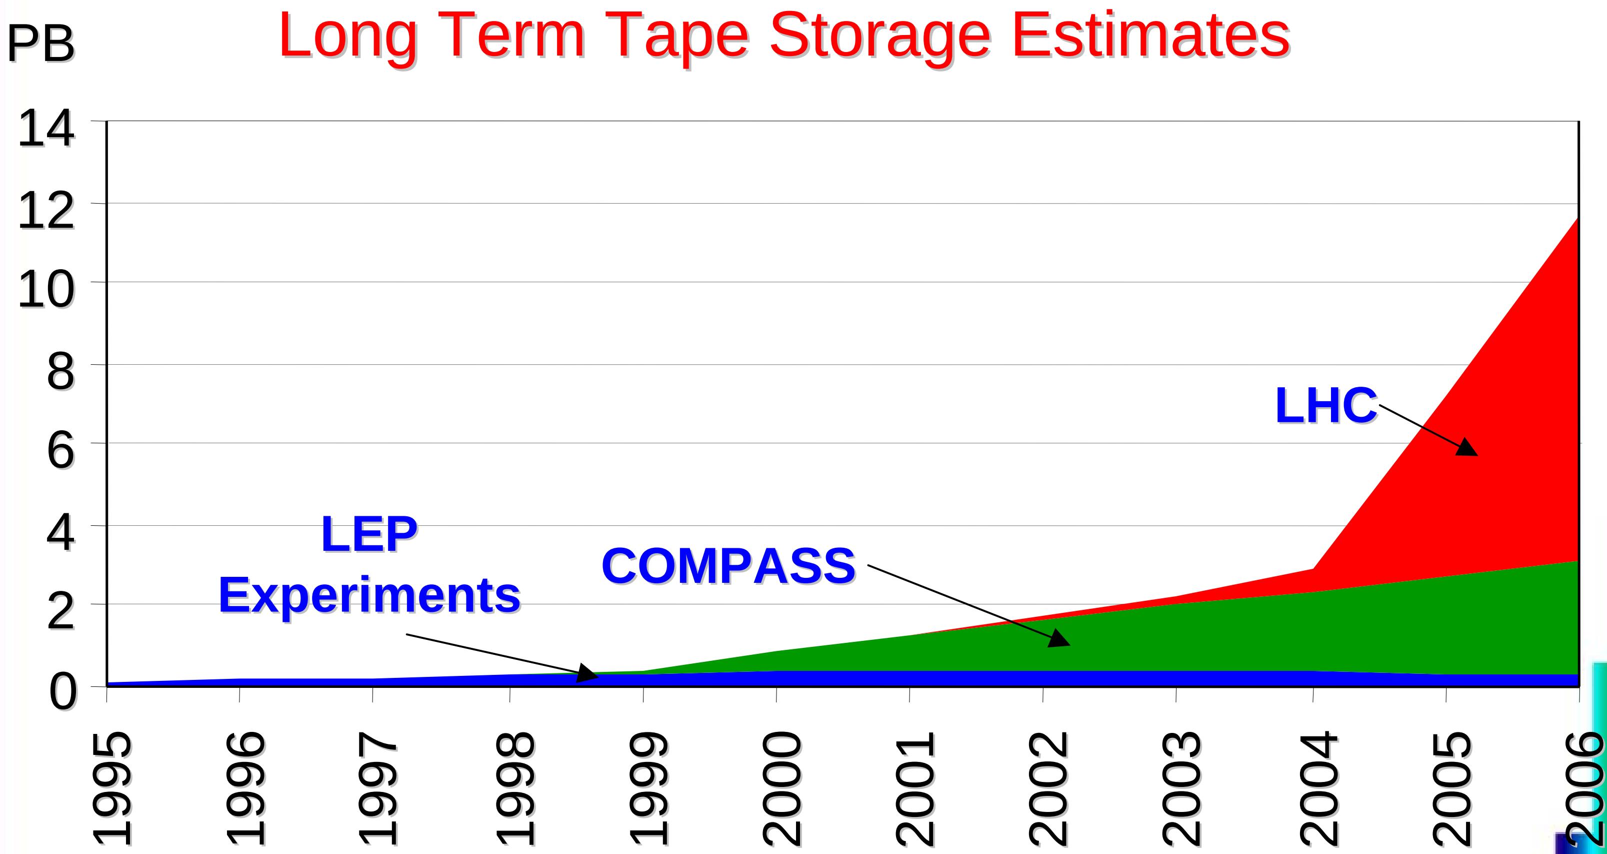

LHC Data Volumes

Data Category	Annual	Total
RAW	1-3PB	10-30PB
Event Summary Data - ESD	100-500TB	1-5PB
Analysis Object Data - AOD	10TB	100TB
TAG	1TB	10TB
Total per experiment	~4PB	~40PB
Grand totals (15 years)	~16PB	~250PB

IBM RAMAC - 1956

- Stored 5 million characters on 50 24 inch disks
- Recording surface painted with same paint as Golden Gate!
- Disk evolution should allow 100TB – 1PB disks towards end of LHC era

Where's the limit?

- Physical limits make prediction beyond 100x today's densities hard
- Future types of storage, e.g. holographic, may provide road ahead

➢ But is there a market for such enormous disks???

- Particularly a commodity market,
  - i.e. your PC

Storage Needs

- Extrapolating from today's reality into future always dangerous
  - T.J.Watson Jr., Ken Olsen, ...
- Will tomorrow's humans record everything that they ever see?
  - From Jim Gray:
    - *1-10GB e-mail, PDF, PPT,*
    - *10-50GB in mpeg, jpeg, ...*
    - *1TB+ voice + video*
    - *Video can drive this towards 1PB*
  - In other words, **1PB of personal data**...

IBM Millipede...

- “The system can store 400 gigabytes per square inch. A prototype, measuring just 3mm square, stores just under 1 gigabyte of data.”
- “in five to 10 years the world may see devices the size of a dime that are capable of storing a terabit of data, which is 125 gigabytes, or 1 trillion bits”
- Rumours that IBM sold its disk business to Hitachi due to Millipede...

Millipede cont.

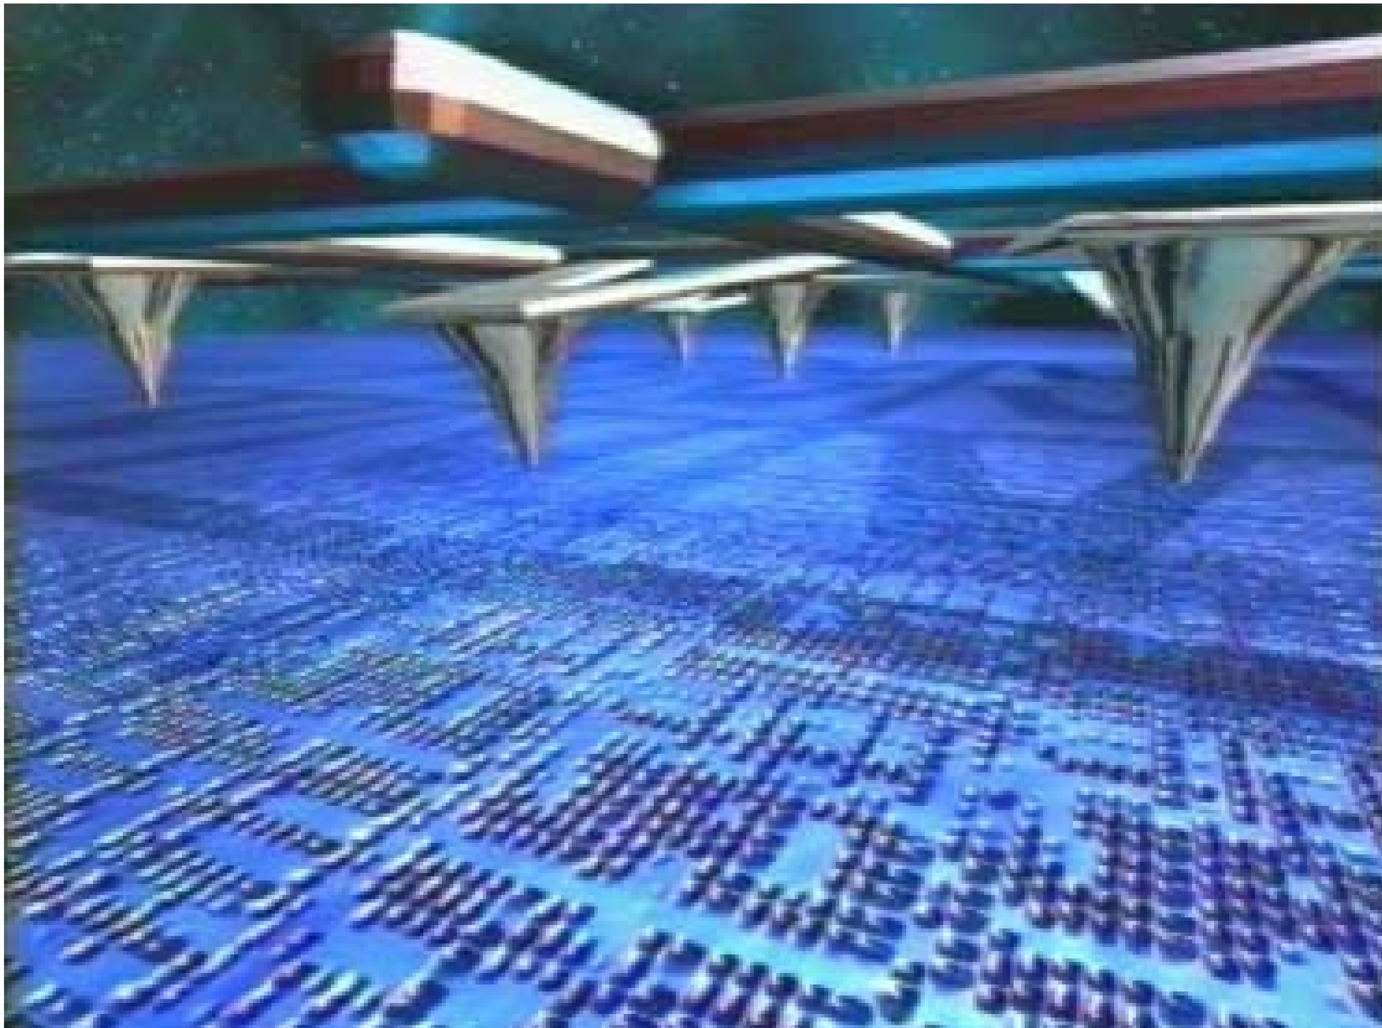

Storage - Predictions

- The personal petabyte

Database Predictions

Databases & HEP

- 1995 on:
  - Distributed Object Database for all data (meta-data, event data, ...)
- Current thinking:
  - Metadata in a database
  - Bulk data in flat files
- LCG Persistency Framework (POOL)
- On-going work with ORDBMS
  - CHORUS, COMPASS, HARP, ...

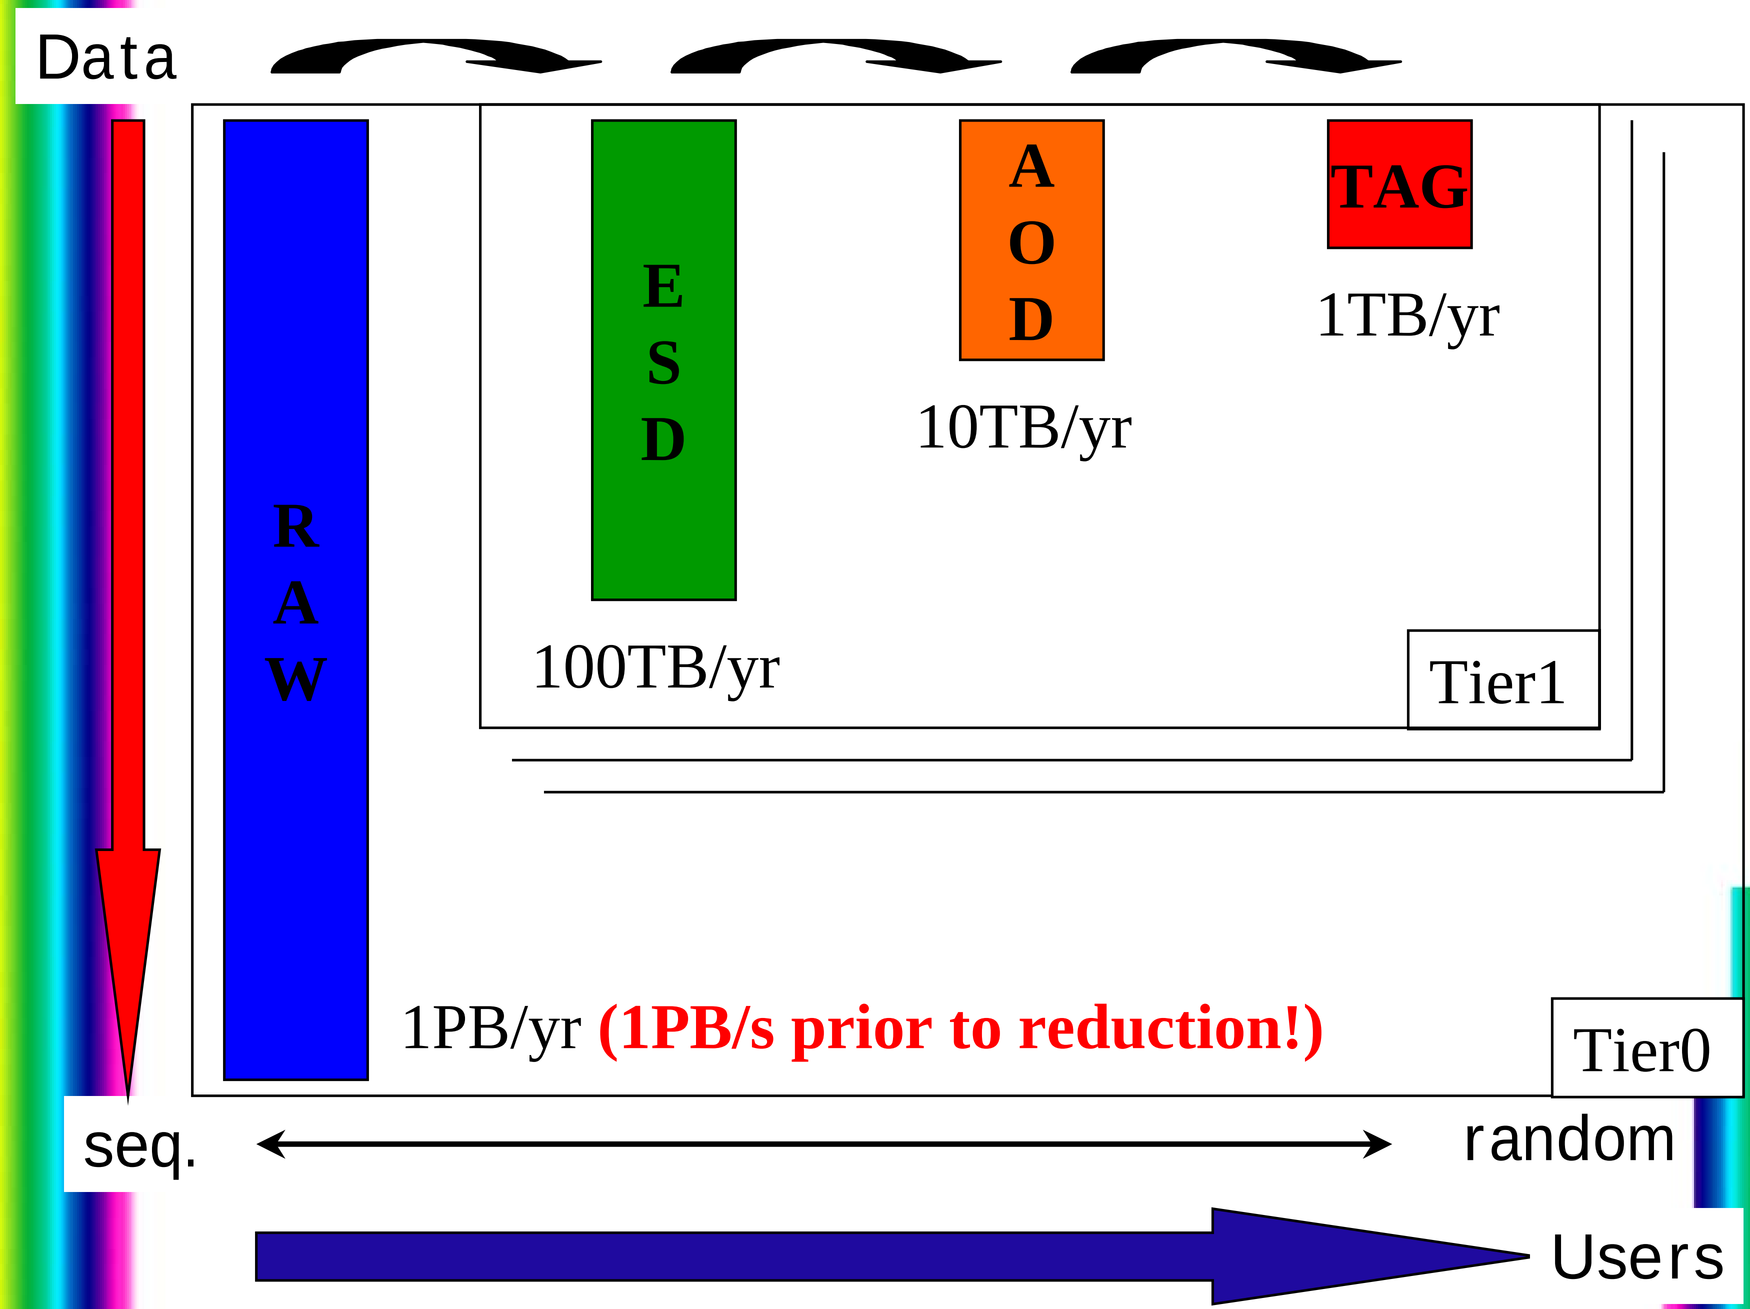
Data
R
A
W
E
S
D
A
O
D
TAG
1TB/yr
10TB/yr
100TB/yr
Tier1
1PB/yr (1PB/s prior to reduction!)
Tier0
seq.
random
Users

Database Predictions

- VLDB: yotabytes by 2020
  - 1,000,000,000 PB

- IBM "Global Technology Outlook"
  - zetabytes by 2010
  - 1,000,000 PB

Marking the Millennium

26th International Conference on Very Large Databases
Cairo, Egypt, 10-14 September 2000

Very Large Data Bases

Reality of Databases Today

- Largest known database: 500TB
  - BaBar experiment at SLAC
- Many databases in 1-10TB range
  - “Management limit” - Jim Gray
- Vendors targetting PB in immediate future

BABAR

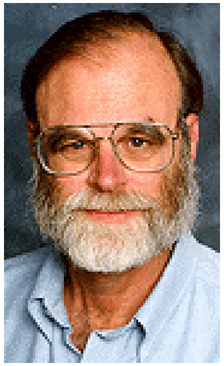

CPU Predictions

Super-Moore's Law

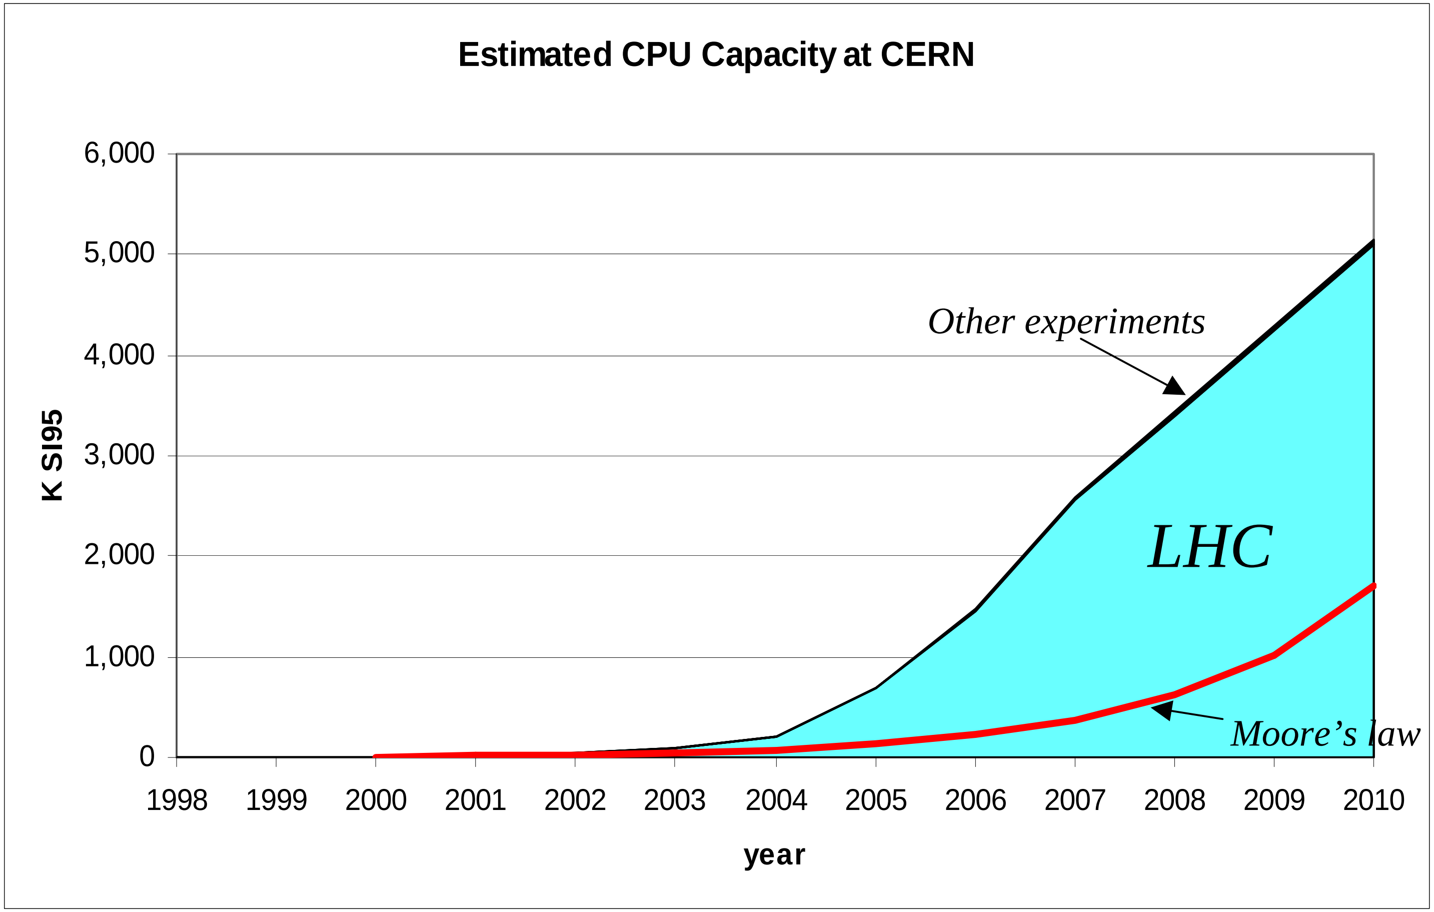

Itanium® Processor Family

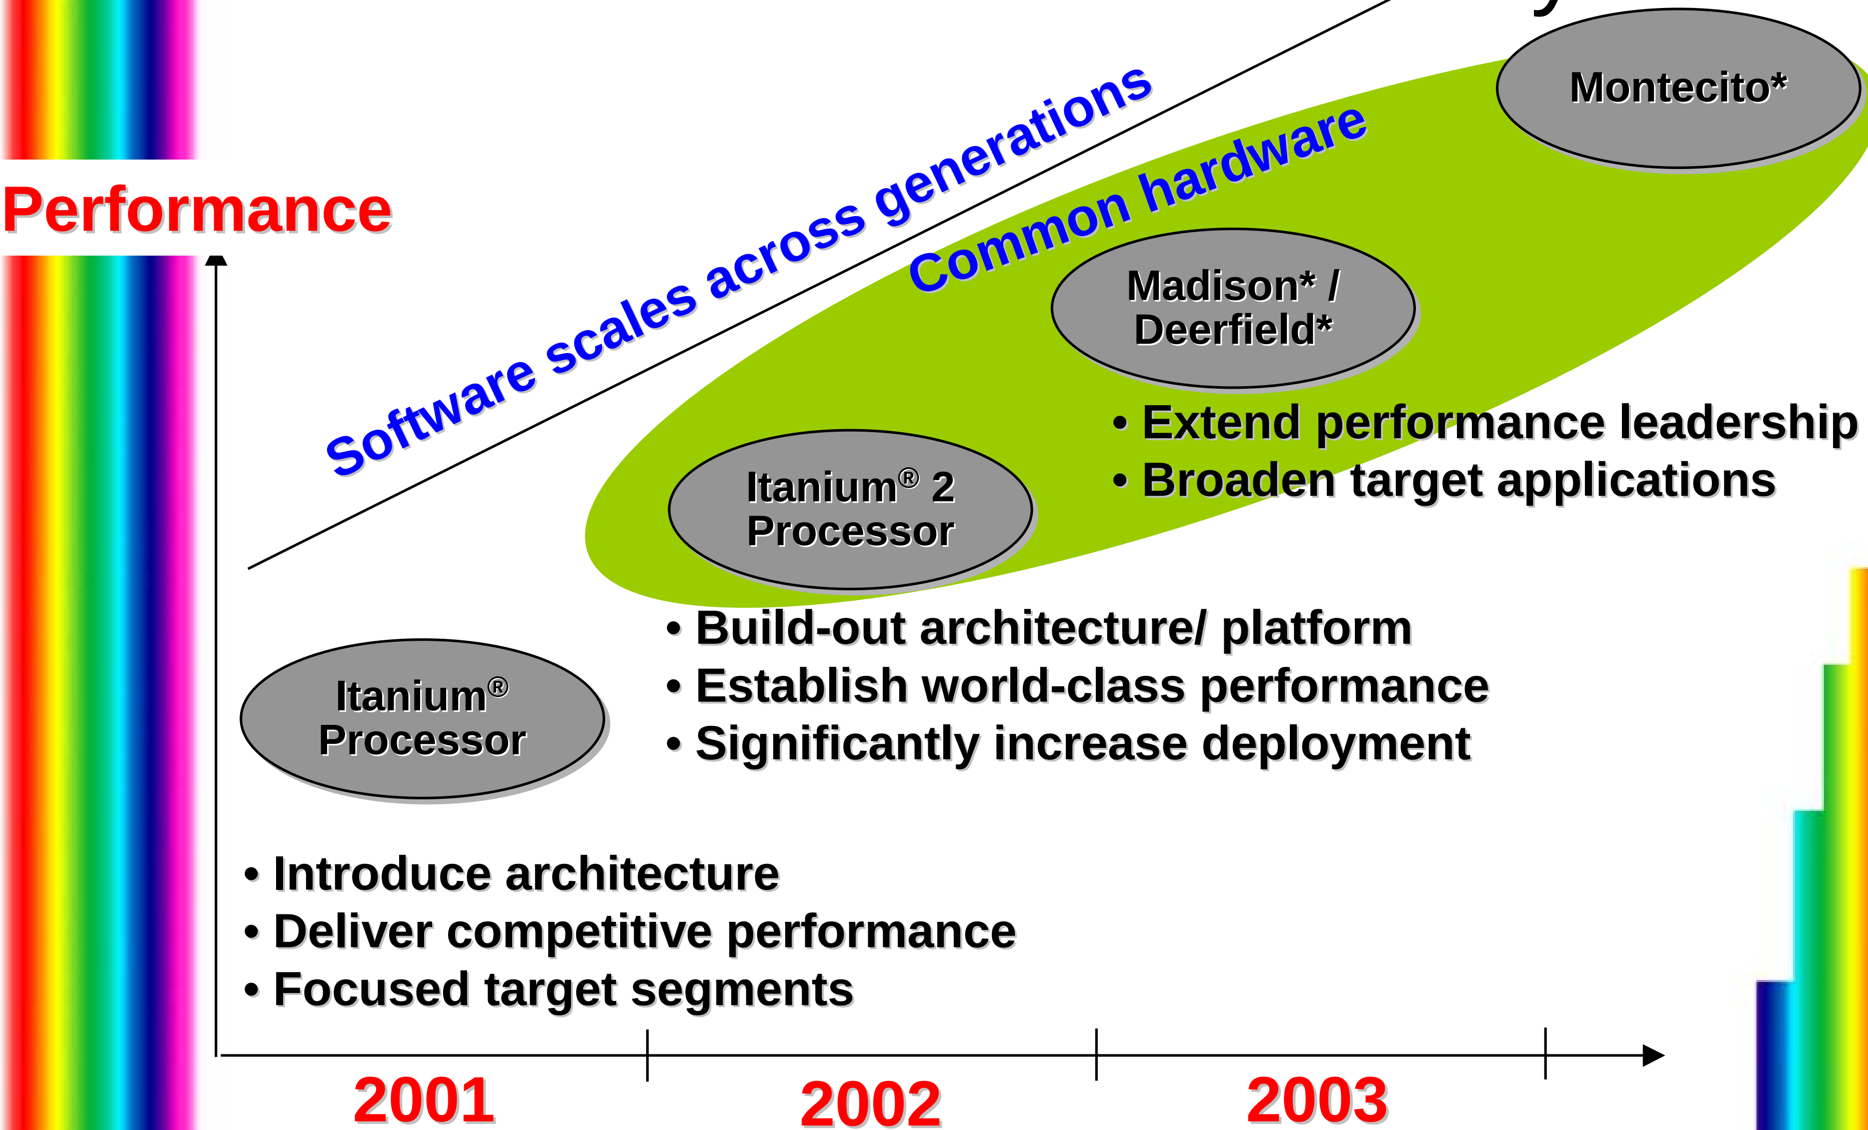

* Indicate Intel processor codenames. All products, dates and figures are preliminary, for planning purposes only, and subject to change without notice.

Grid

S...

- “A … n
  w … er you
  di …
  re …
  ur …

Les …

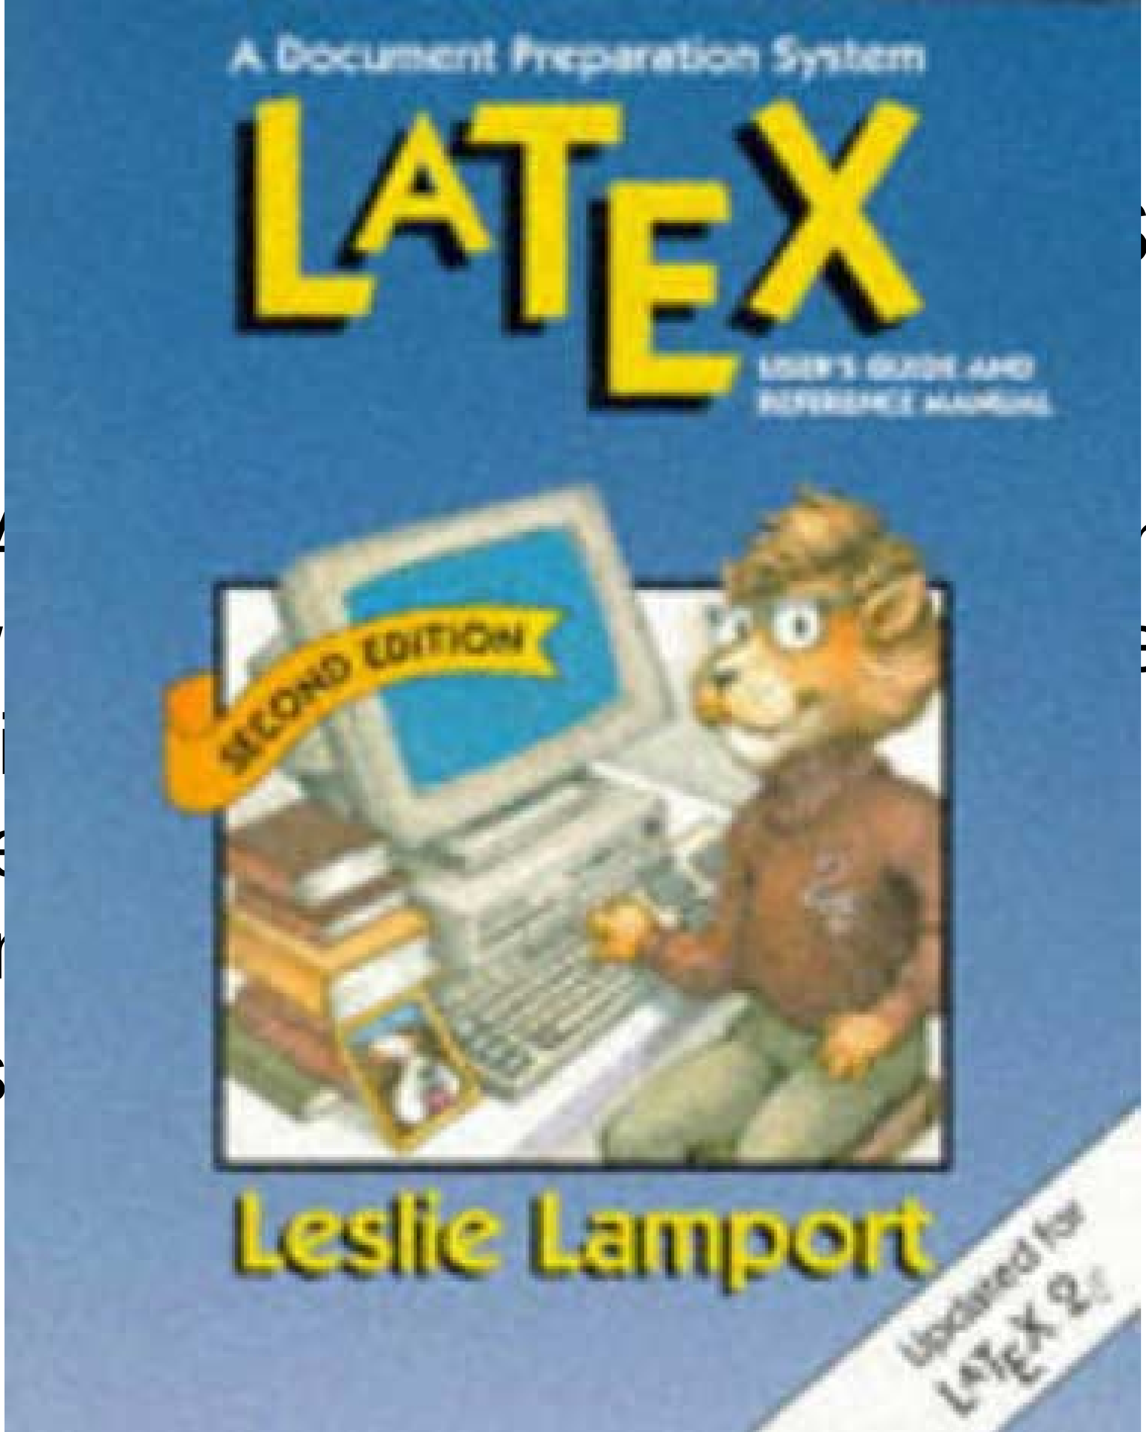

I g

- "If obably woul ts about to b
- "We ns of com
  1. M
  2. Cl
  3. In
- "Th cture for com years"

The Grid

- Overview – see DG's introductory talks
- Detail – see Tony Hey's talk on August 21
  - eBusiness, eScience & the Grid
- CERN & the Grid
  - Many projects, specifically:
  - EU Data Grid (EDG)
  - LHC Computing Grid (LCG)

The Grid vision

- Flexible, secure, coordinated resource sharing among dynamic collections of individuals, institutions, and resource
  - From “The Anatomy of the Grid: Enabling Scalable Virtual Organizations”
- Enable communities (“virtual organizations”) to share geographically distributed resources as they pursue common goals -- *assuming the absence* of…
  - central location,
  - central control,
  - omniscience,
  - existing trust relationship

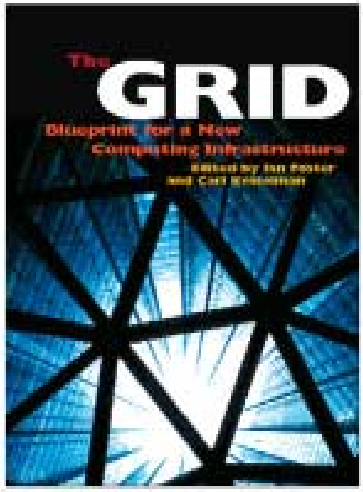

the globus project
www.globus.org

Grids: Elements of the Problem

- Resource sharing
  - Computers, storage, sensors, networks, …
  - Sharing always conditional: issues of trust, policy, negotiation, payment, …
- Coordinated problem solving
  - Beyond client-server: distributed data analysis, computation, collaboration, …
- Dynamic, multi-institutional virtual orgs
  - Community overlays on classic org structures
  - Large or small, static or dynamic

Grid R&D Projects

= EDG

Many national, regional Grid projects -- GridPP(UK), INFN-grid(I), NorduGrid, Dutch Grid, …

US projects

European projects

EDG Interfaces

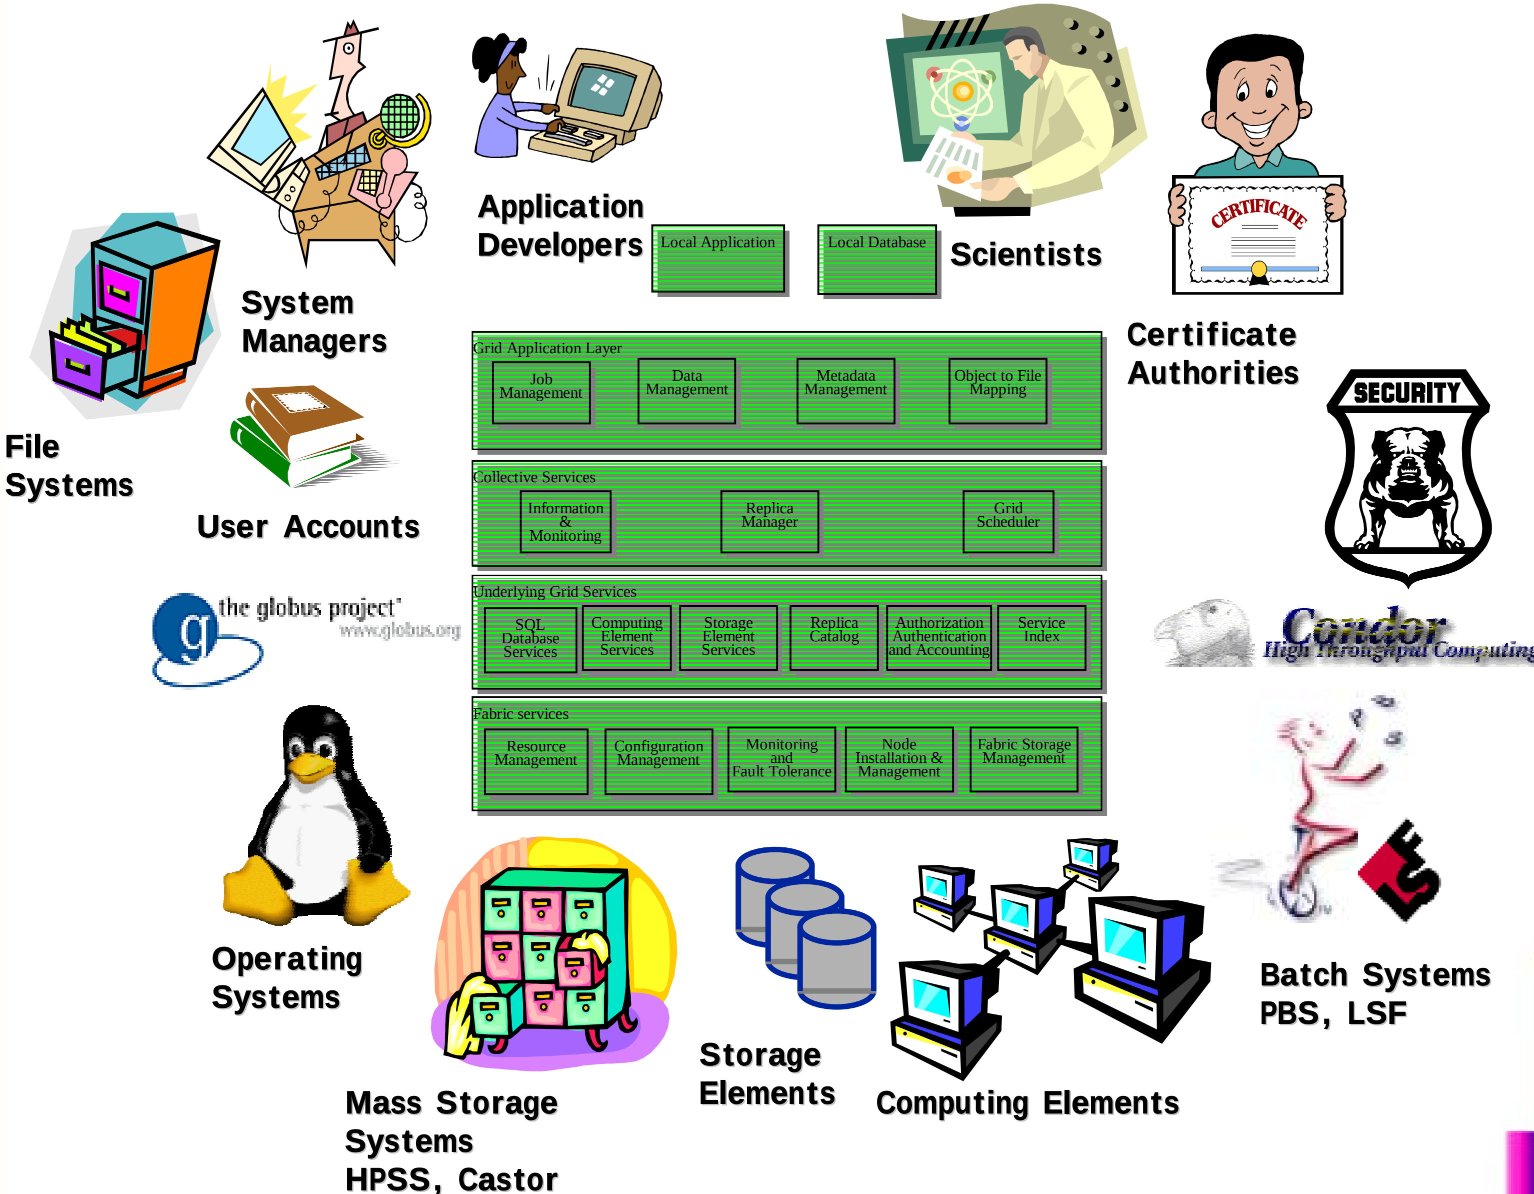

Biomedical applications

- Data mining on genomic databases (exponential growth)
- Indexing of medical databases (Tb/hospital/year)
- Collaborative framework for large scale experiments (e.g. epidemiological studies)
- Parallel processing for
  - Databases analysis
  - Complex 3D modelling

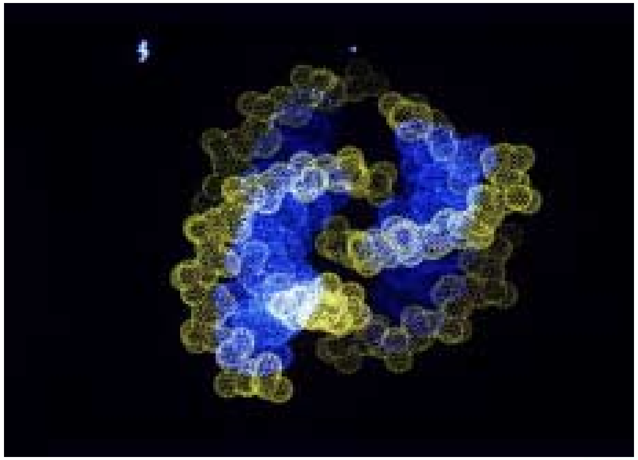

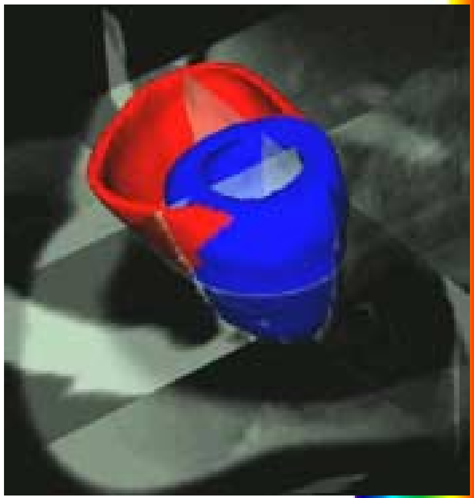

Earth Observations

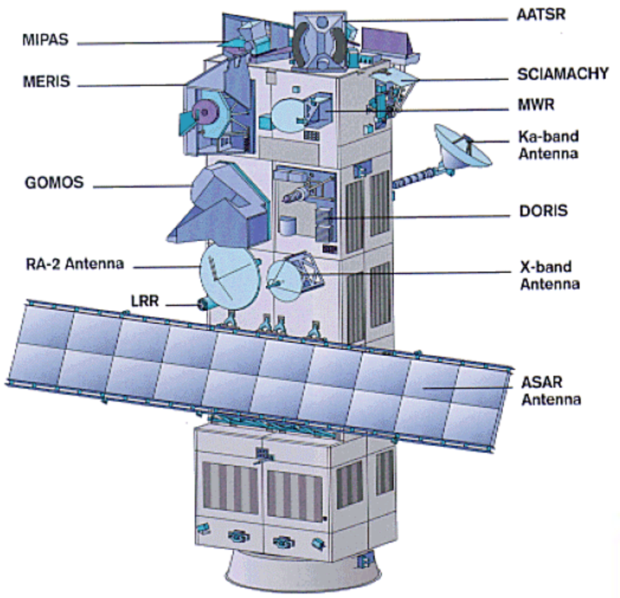

ESA missions:

- **about 100 GB of data per day (ERS 1/2)**
- **500 GB for the next ENVISAT mission (launched March 1st)**

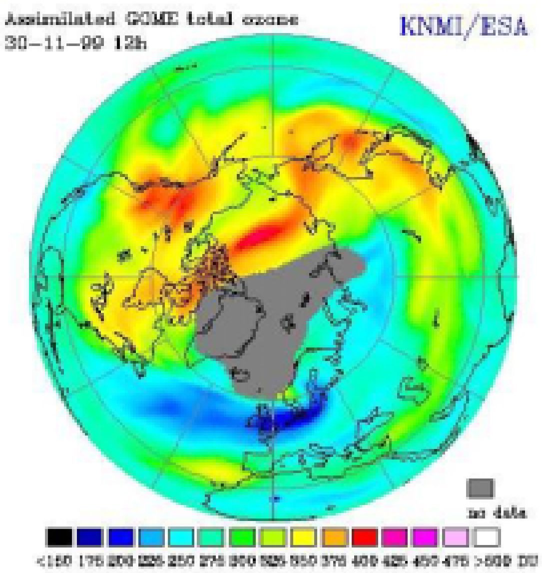

EO requirements for the Grid:

- **enhance the ability to access high level products**
- **allow reprocessing of large historical archives**
- **improve Earth science complex applications (data fusion, data mining, modelling ...)**

Grids & Industry

- Strong push from major vendors, including IBM and others
  - e.g. Sun, Microsoft, ...

- Consistent message of Grid as “next generation of Internet”
  - Networking (TCP/IP)
  - Communications (e-mail)
  - Information (World Wide Web)
  - Computing (Grid)

The Internet as a Computing Platform

Key Challenges / Initiatives

- Building an open infrastructure
  - Grid Computing Protocols
- Managing the infrastructure
  - Autonomic Computing
- Accessing the infrastructure
  - "Utility" Computing

Autonomic Computing

- Self-optimizing
- Self-configuring
- Self-healing
- Self-protecting

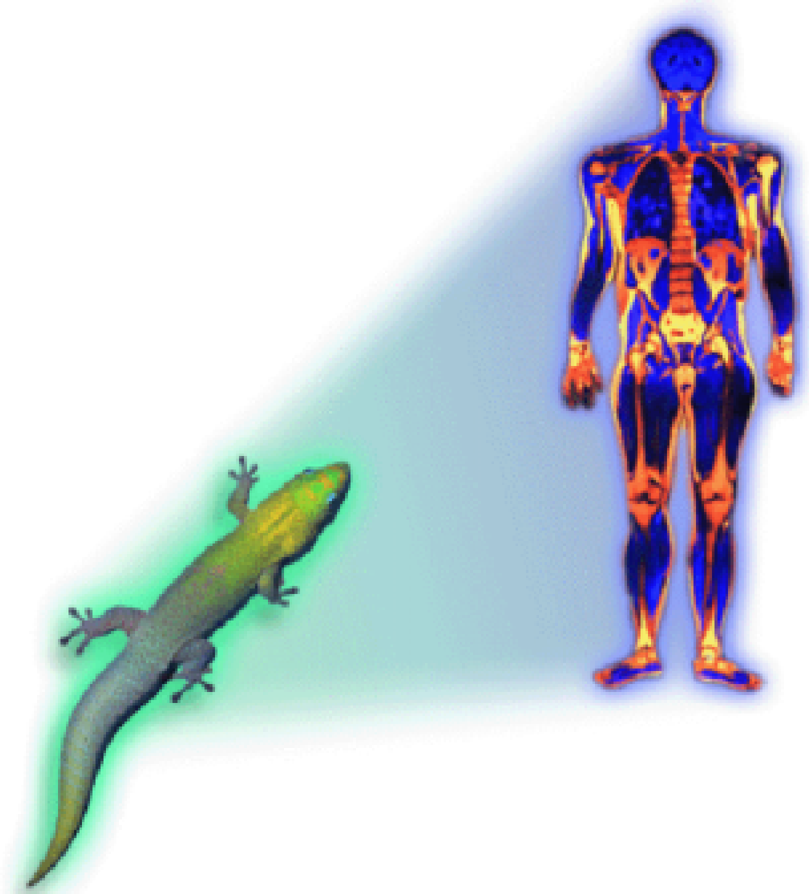

"Utility" Computing Examples

Canadian Tire

Dow Chemical

Memphis Schools

Saks Incorporated

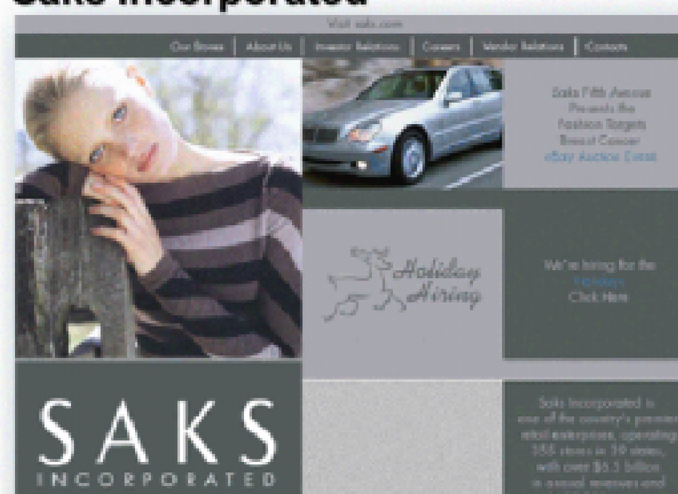

Computing Predictions

Wearable Computers

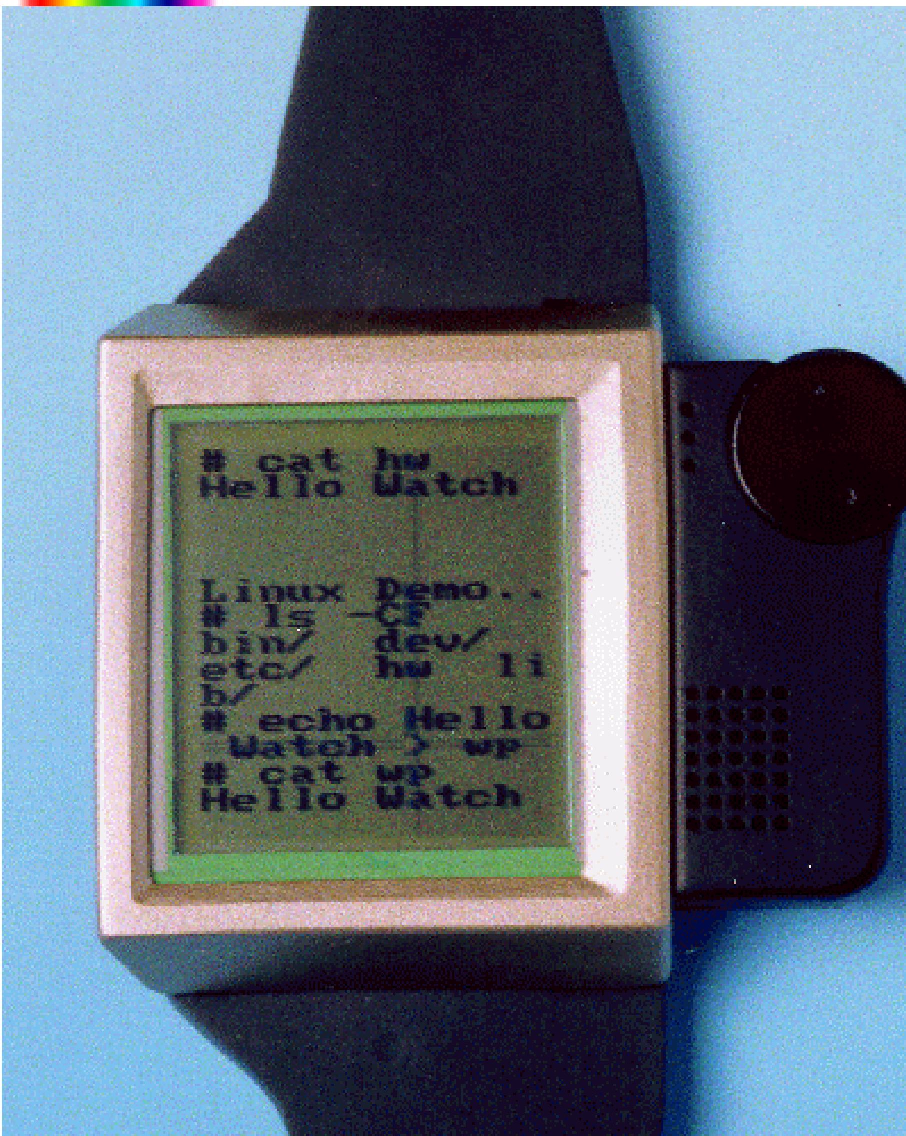


Copyright © 1989 by Paramount Pictures. All Rights Reserved.

Smart Dust

- Develop complete sensor / communication system into 1 mm^3
  - "Grain of sand" also mentioned...
- Potential applications:
  - Virtual keyboard
  - Inventory control
  - Product quality monitoring
  - Smart office spaces

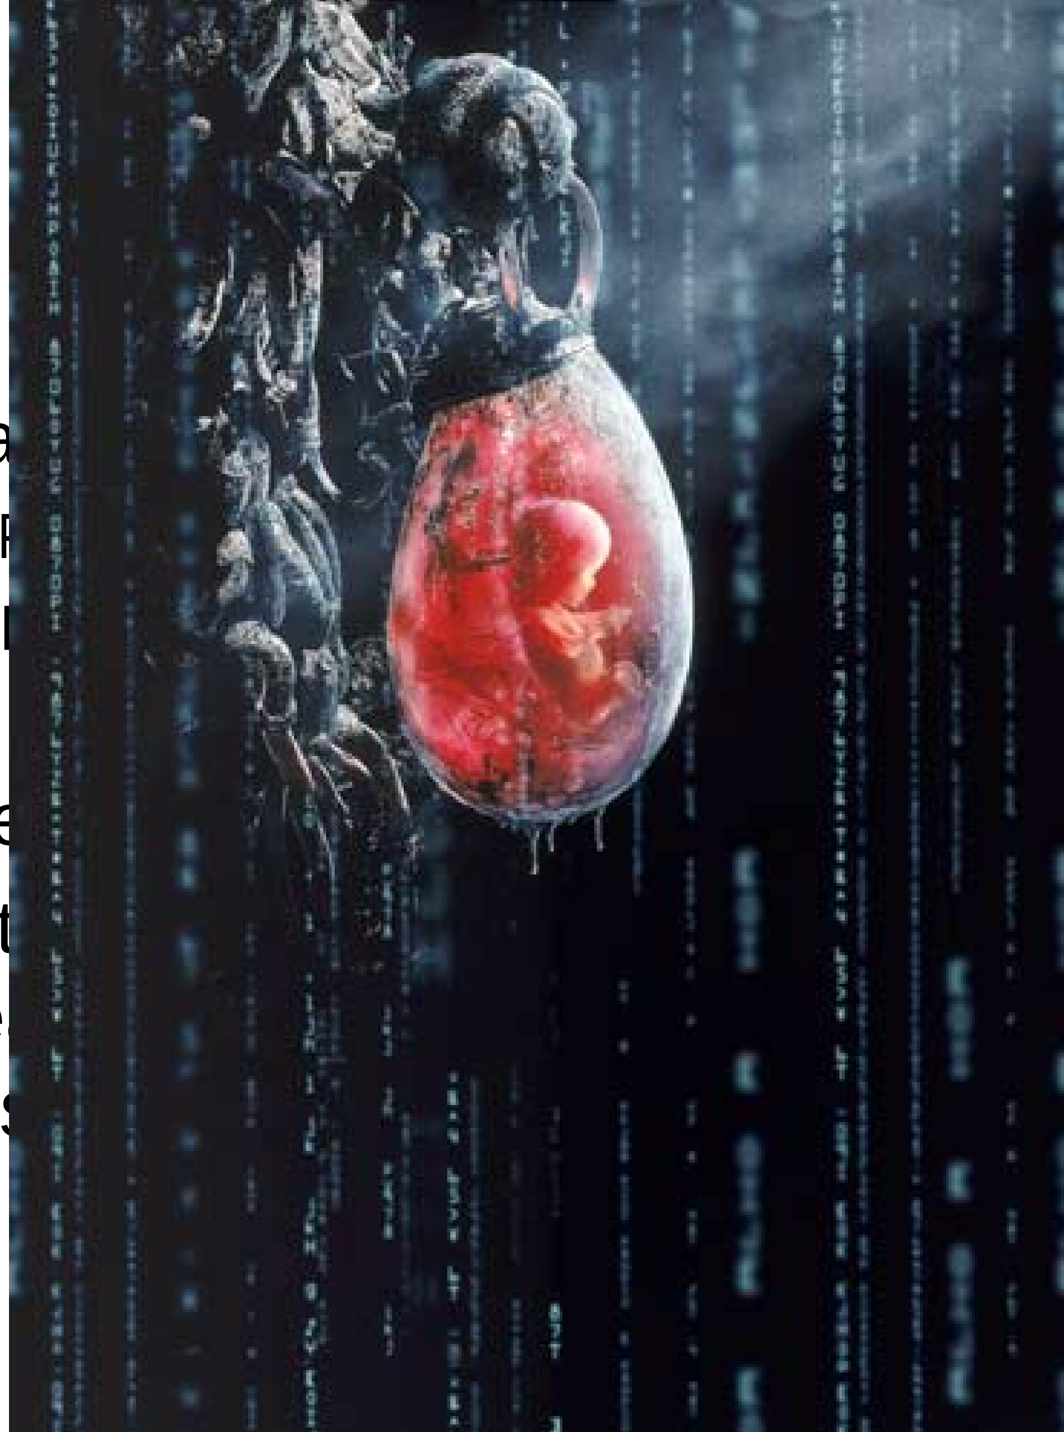

- Ma
  - F …
  - I ours
- IB nduction ove
- Alt ole = we
  - S cars

Smart Dust again

- “Scavenging power from sunlight, vibration, thermal gradients, and background RF, sensors motes will be immortal, completely self contained, single chip computers with sensing, communication, and power supply built in.
- Entirely solid state, and with no natural decay processes, they may well survive the human race. Descendants of dolphins may mine them from arctic ice and marvel at the extinct technology.”

The last 100 years...

Population	4
Horses	1.1
Forest area	0.8
Blue whales	0.0025 (1/400)
World economy	14
Energy use	13
CO_2 emissions	17
Industrial output	40
Computers	?

Predictions from 1945

- "As we may think"
  - Vannevar Bush
- Describes "memex"
  - A memex is a device in which an individual stores all his books, records, and communications, and which is mechanized so that it may be consulted with exceeding speed and flexibility. It is an enlarged intimate supplement to his memory.
- Used in much the same way as the Web

Predictions from 2000

- In 2010, everything worth more than a few $ will know that its yours...
- A speck of dust on each fingernail will communicate with your computer
- Your house, office and car will be continuously aware of your presence
- Tyres will communicate with the on-board computer if pressure is low, your milk carton will signal if the contents are off...
- In 2020, sensors will monitor all major bodily systems, providing early warning of diseases...

Summary

Summary I

- We've looked at:
  - The birth of IBM,
  - The IBM PC,
  - Unix, then Linux,
  - The Internet, The Web,
  - GUI / mouse, ...

Summary II

Producing high-quality software is:

- Far from easy
- Far from cheap
- Still not a solved problem

Discussion Session

Friday 26th July, 11:15, main amphitheatre

Further Reading

Some Links

http://www.h2g2.com/

http://www.bbc.co.uk/cult/doctorwho/

http://cern.ch/ssl-computing/default.htm

Acknowledgements

Many in IT, CERN and anyone who's put something on the Web

Homework

BACK TO THE FUTURE

End Lecture I I I