Fermilab U.S. DEPARTMENT OF ENERGY Office of Science

Computing Innovations

Oliver Gutsche

Fermilab 50th Anniversary Symposium

7. June 2017

Computing - It's a tool for science!

- It's all about the science!
- Computing
  - It should be easy, fast, enabling
- It should just be there!

The Challenge - It is not that easy!

"The graphics displays
you provide are
inadequate
to
look in the Eye of God!"

Rocky Kolb to Hugh Montgomery, sometime in the 1980's

Advancement Cycle

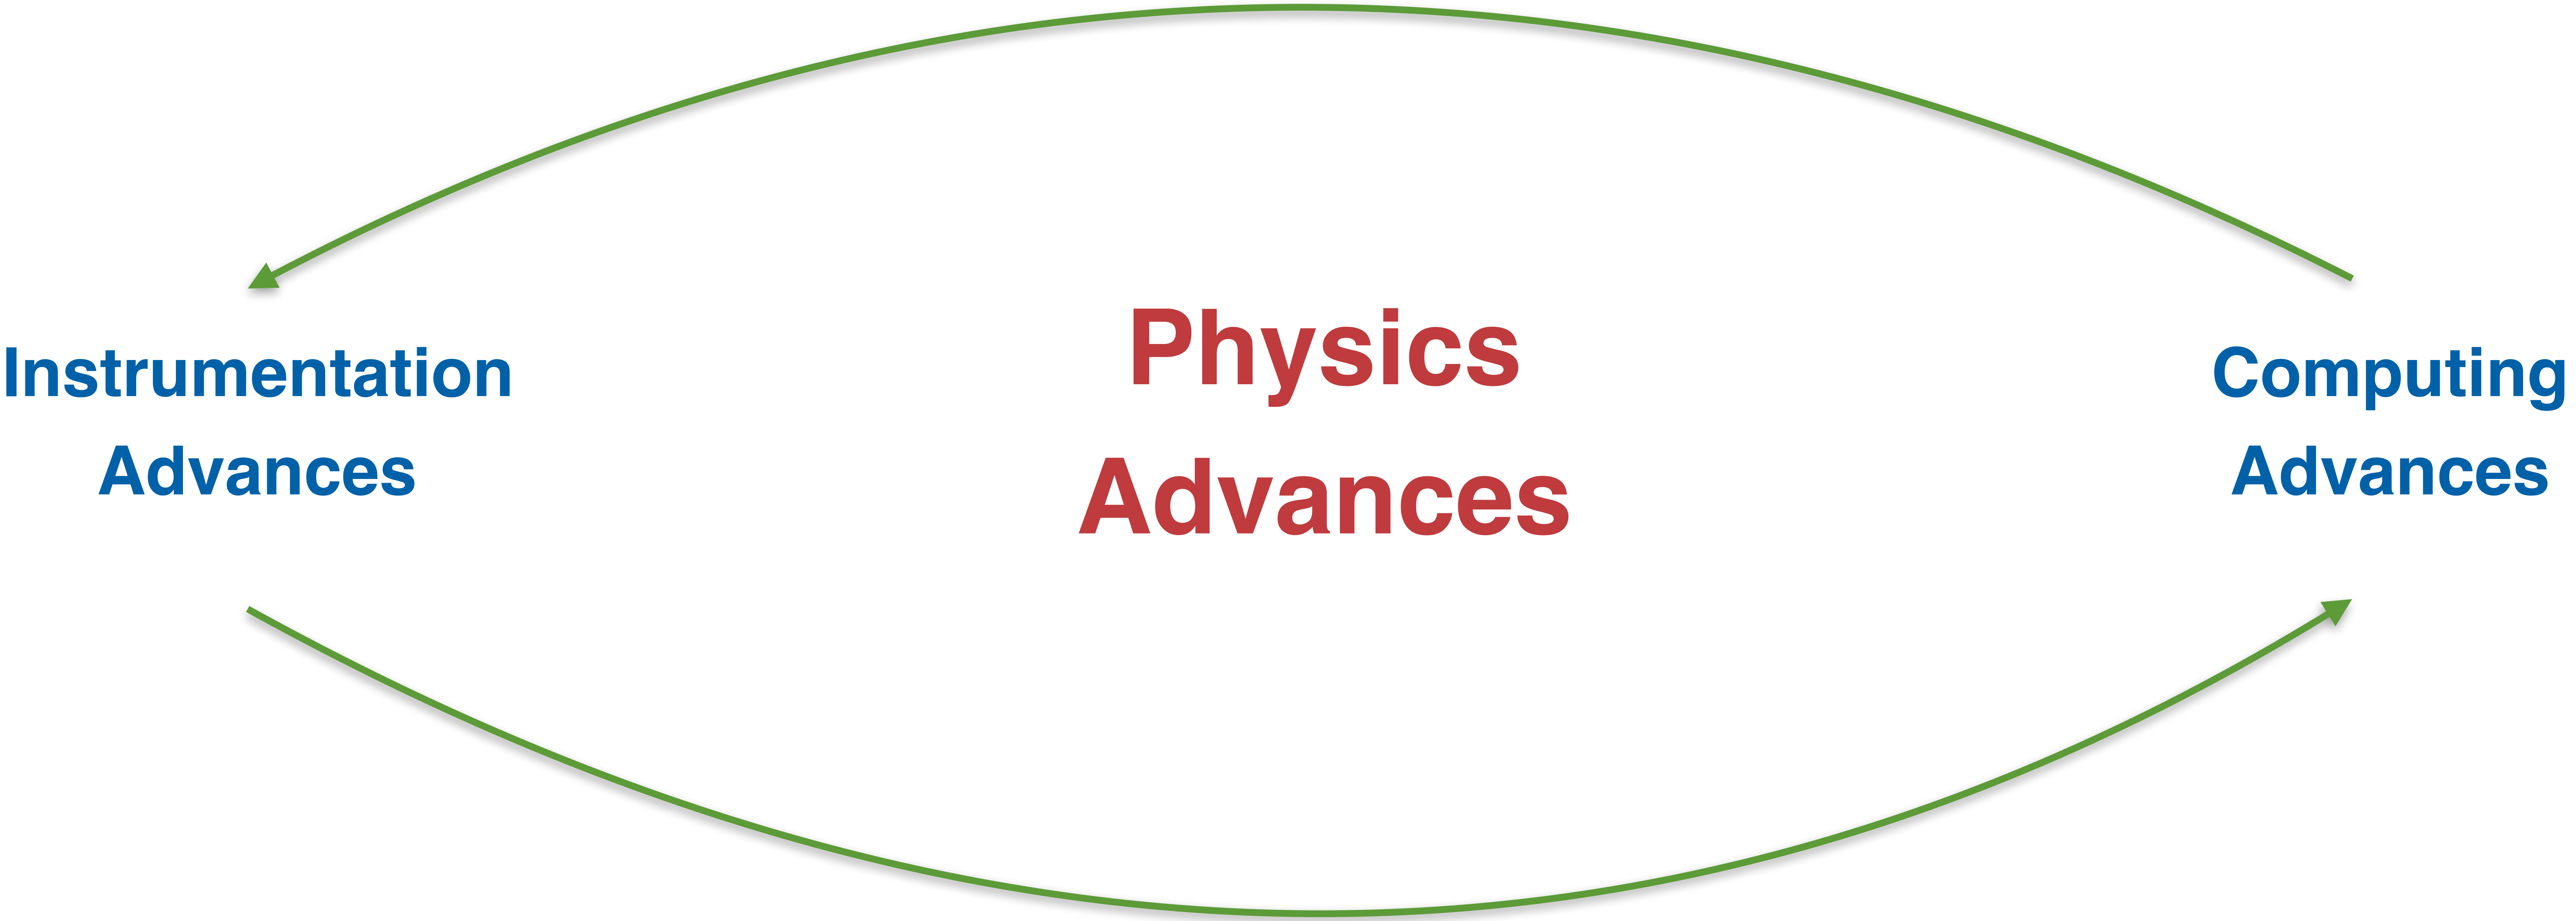

The early days

- No dedicated computing department
- Computing was done offsite, at Argonne National Laboratory and elsewhere
- Computers were primarily used to design accelerators rather than for analysis or simulations of physics.

Fermilab

The first mainframes

1976

Cybers

There is tape and even disk

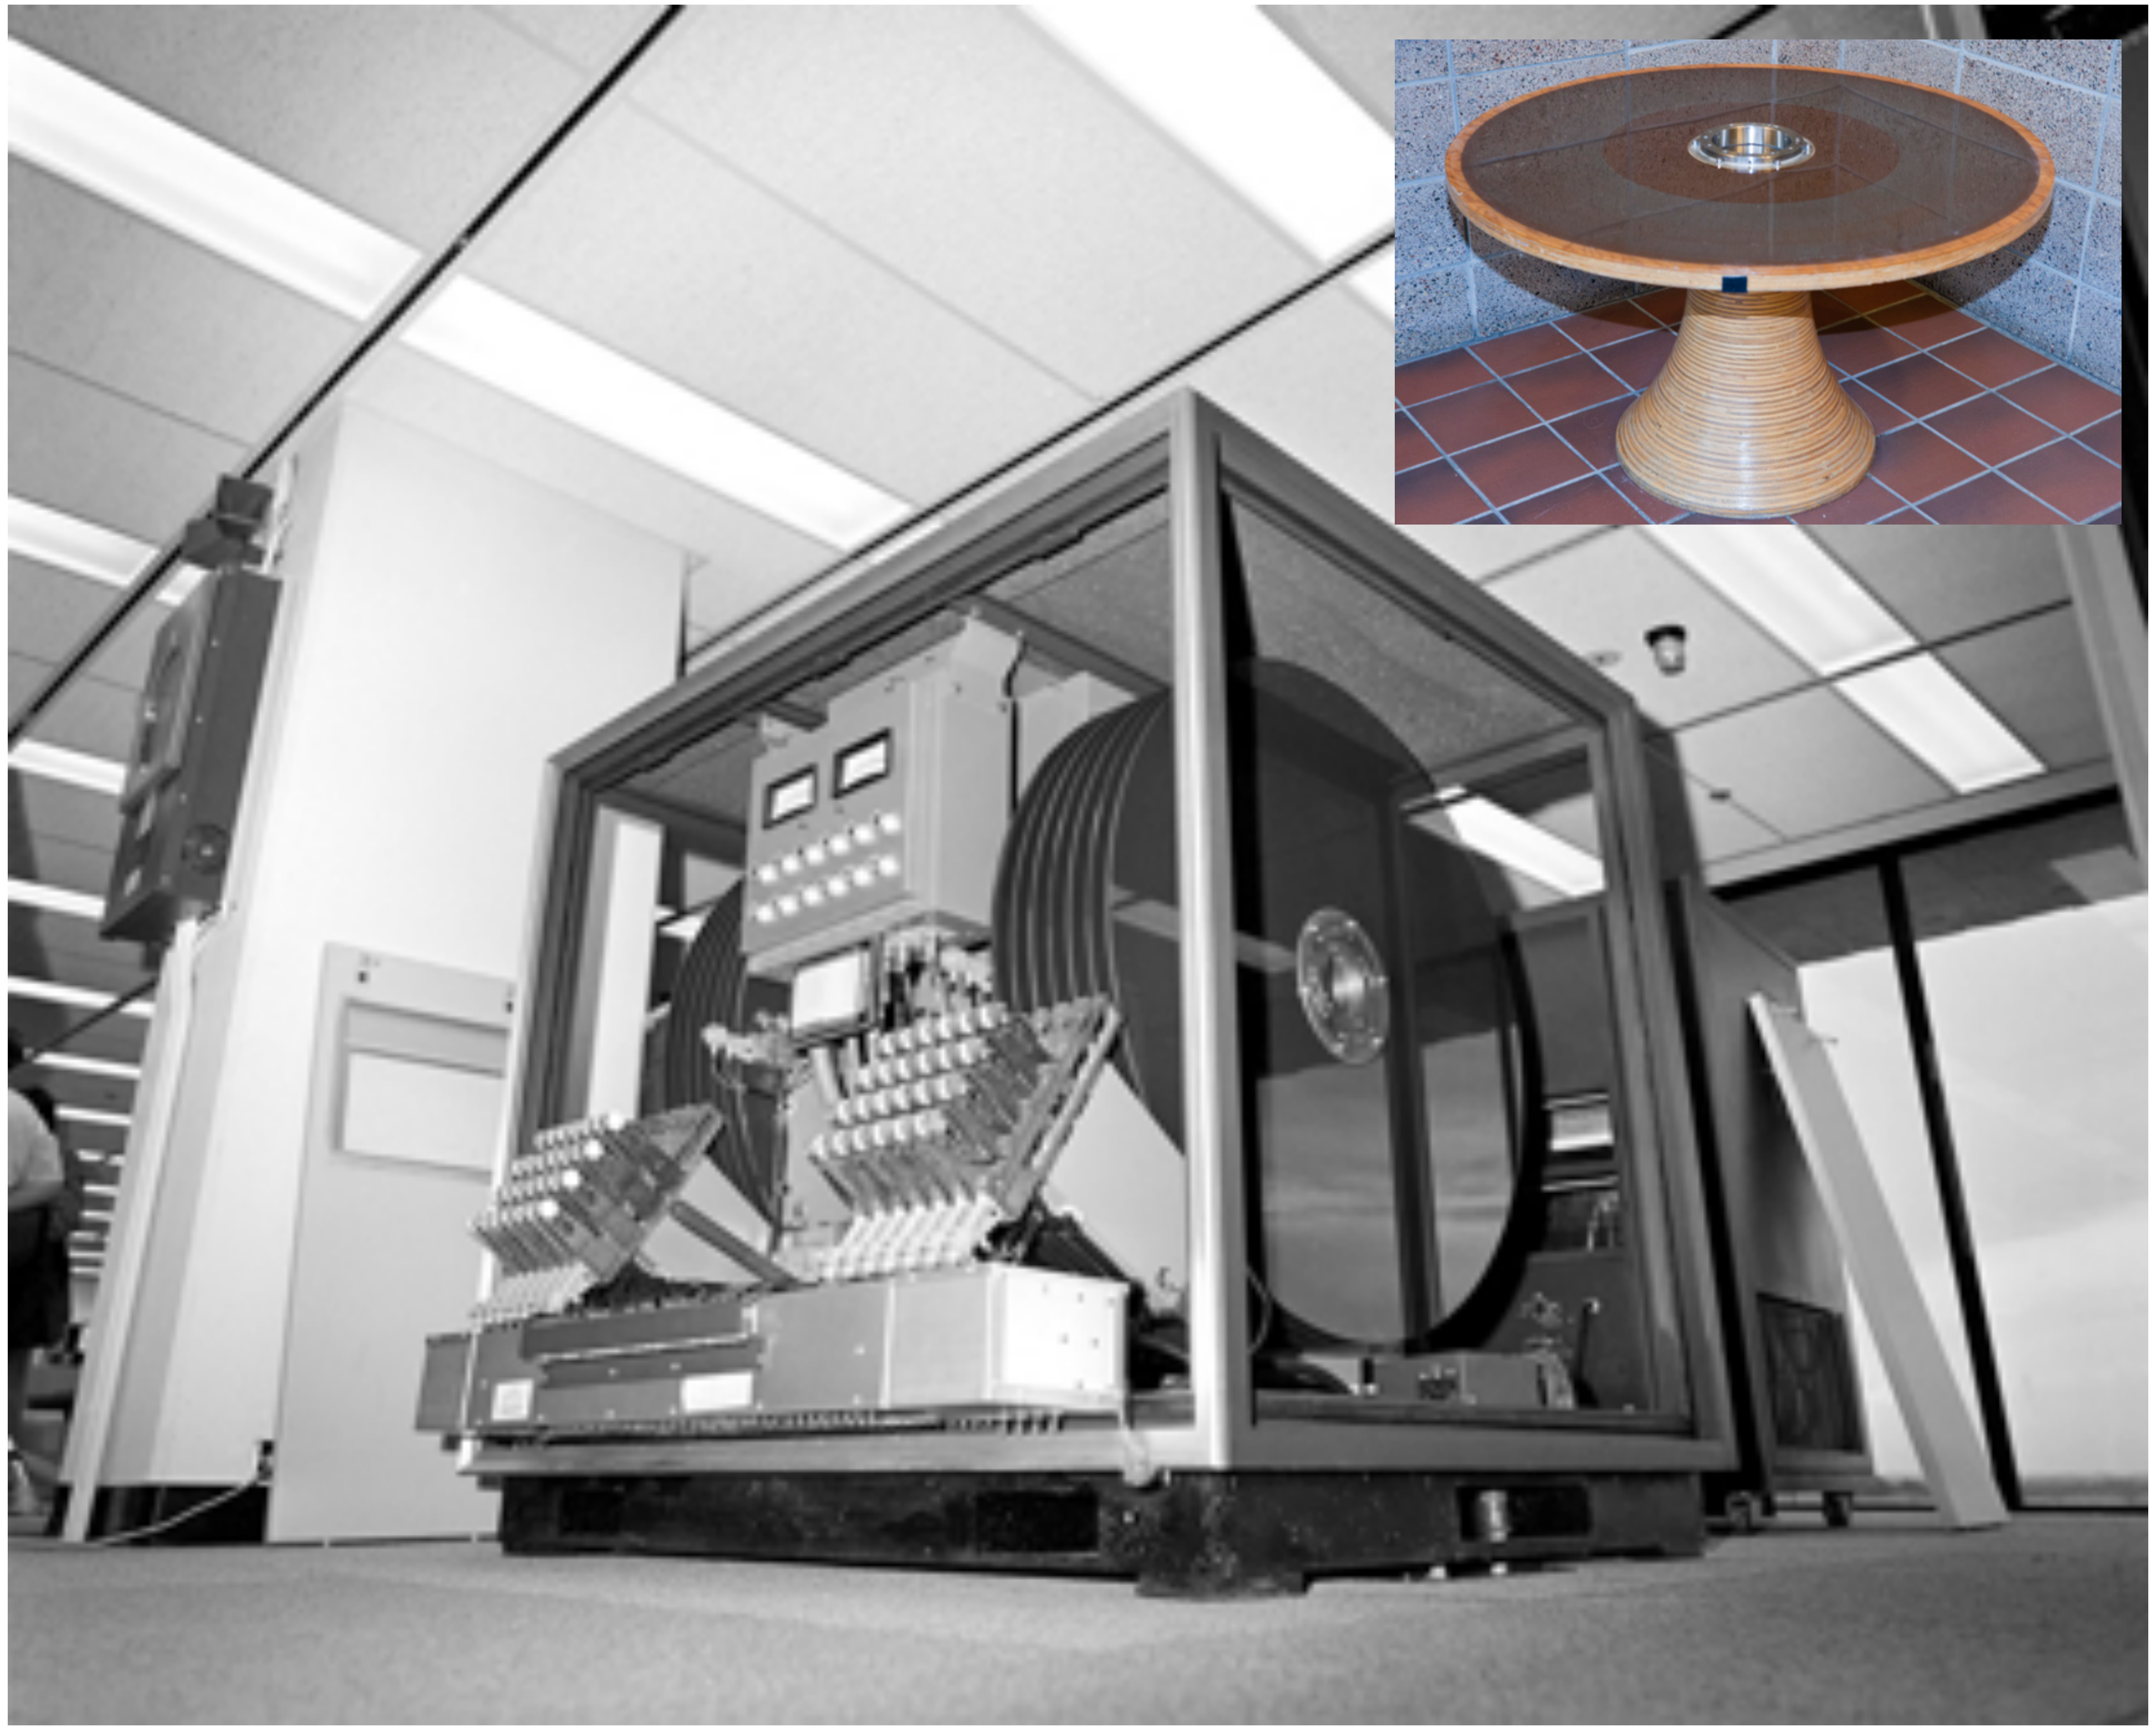

Online innovations drive offline innovations

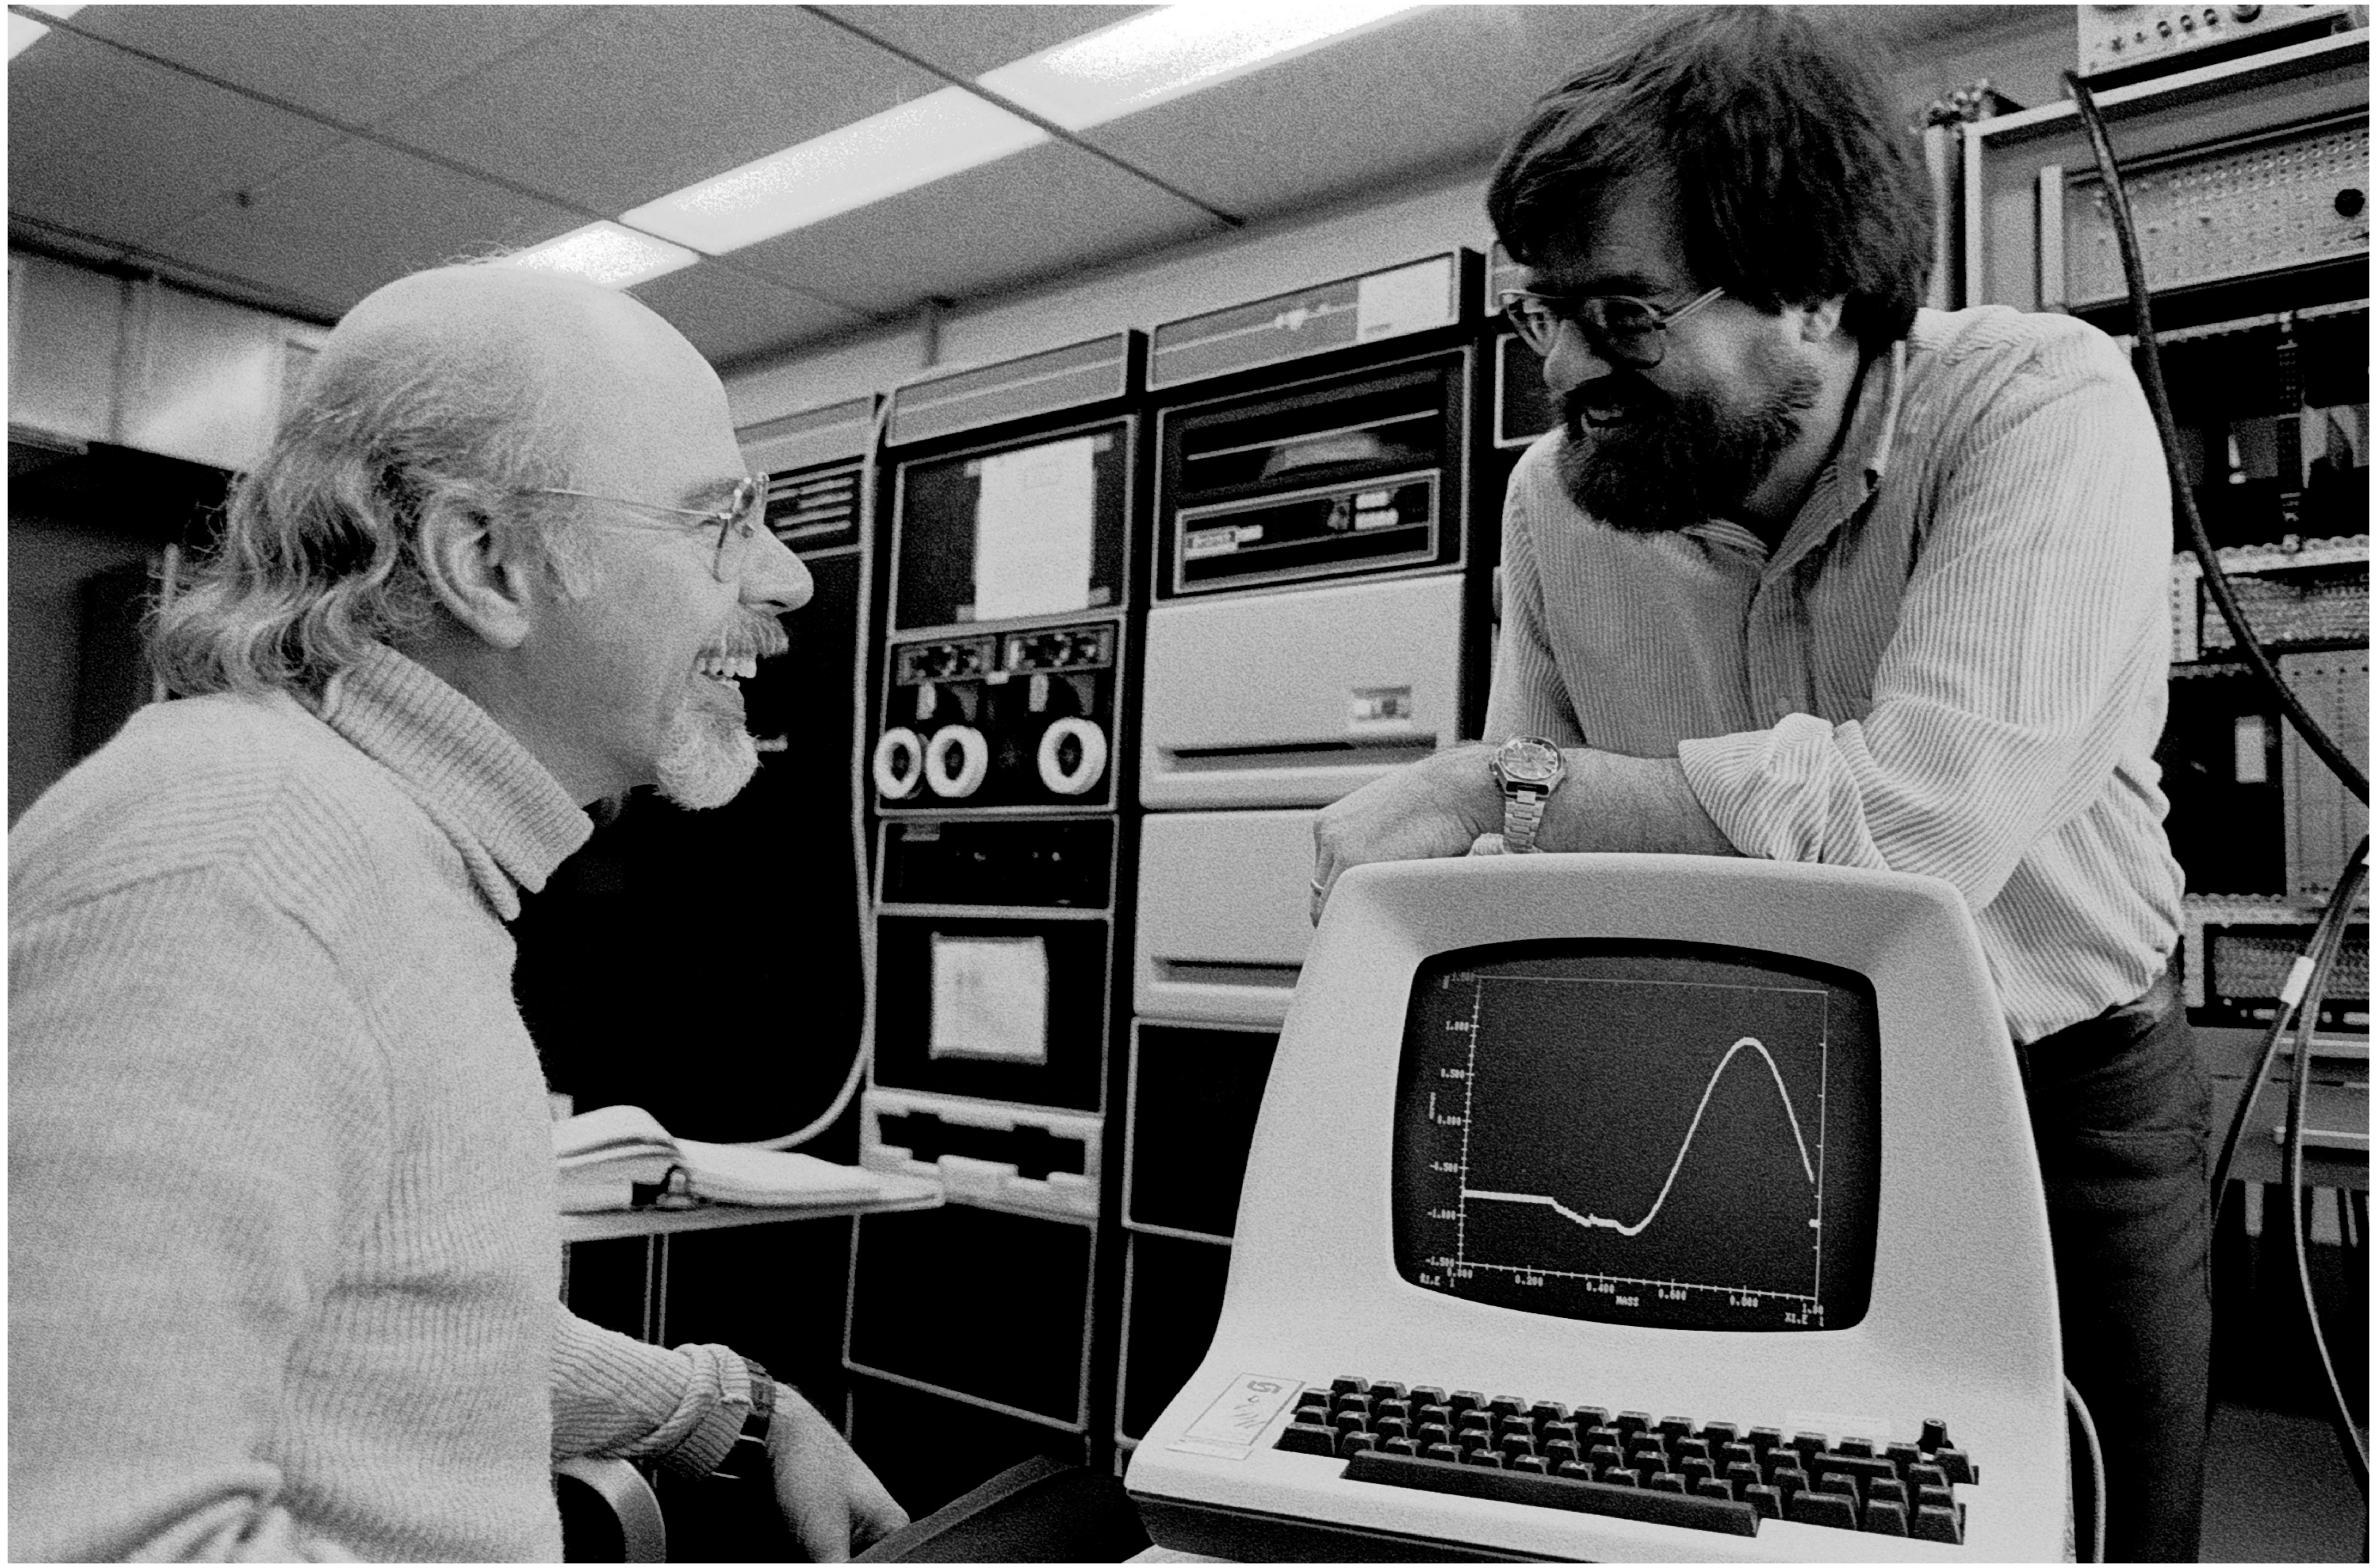

- Data Acquisition (DAQ): First PDP 11

Fermilab

Software development: punch cards

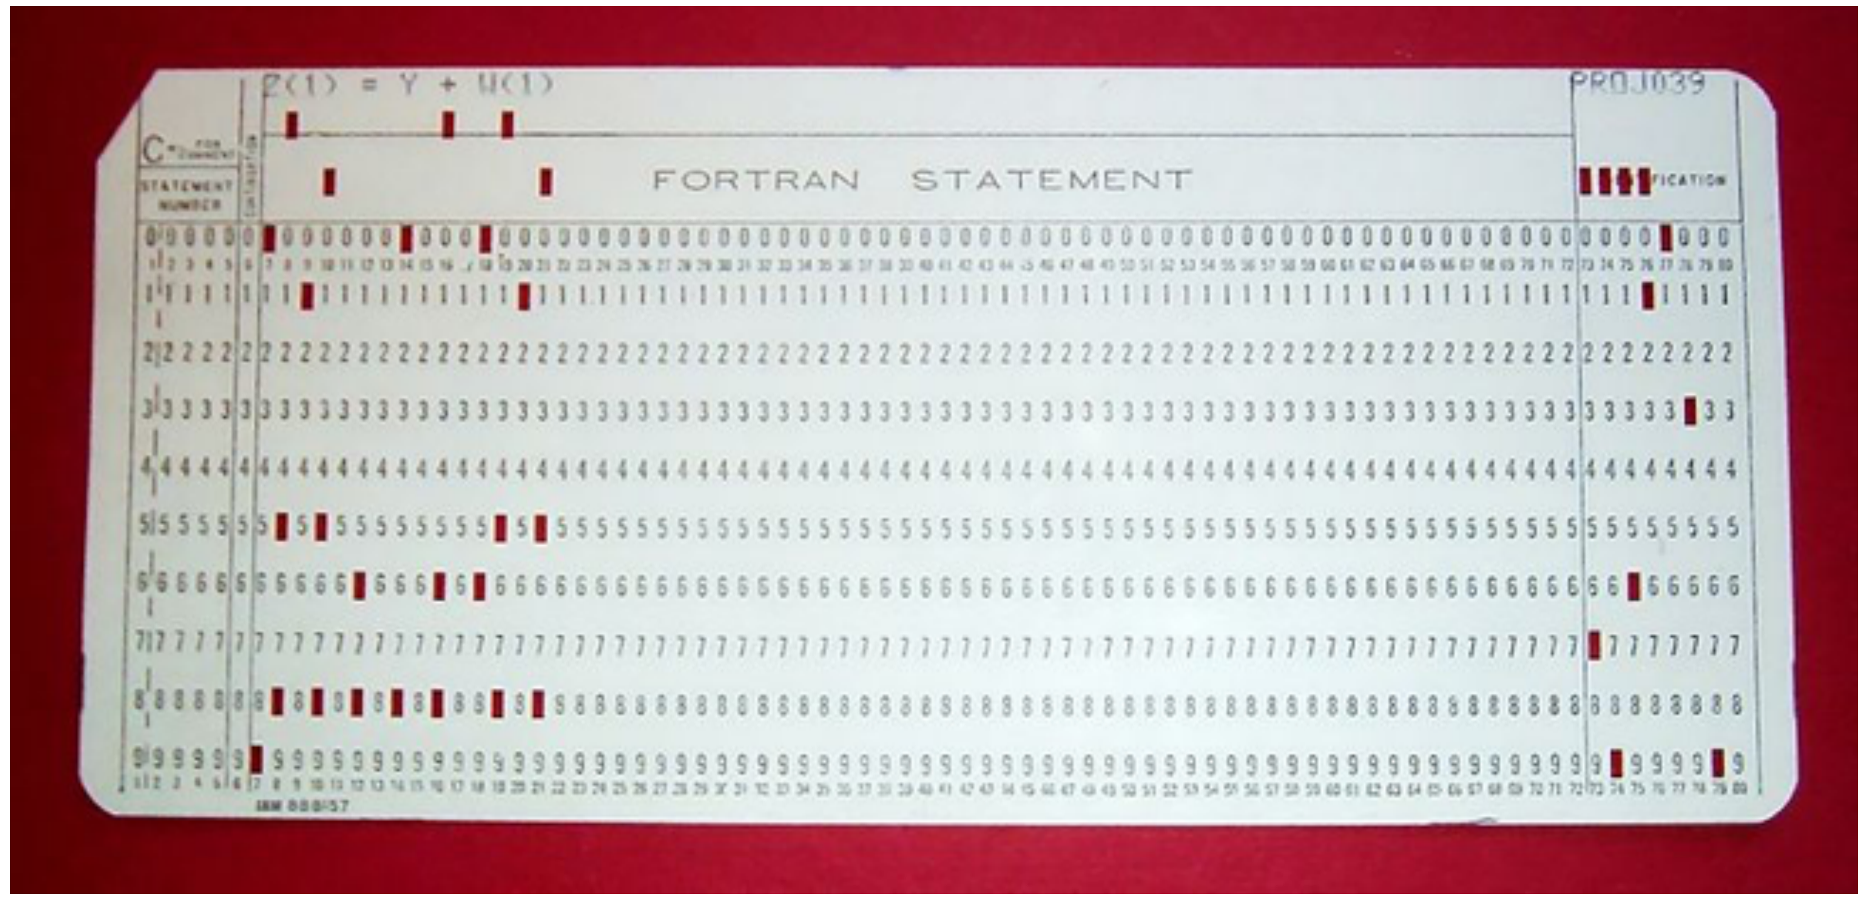

Network: VAX ➜ DECNet ➜ HEPNet

1982

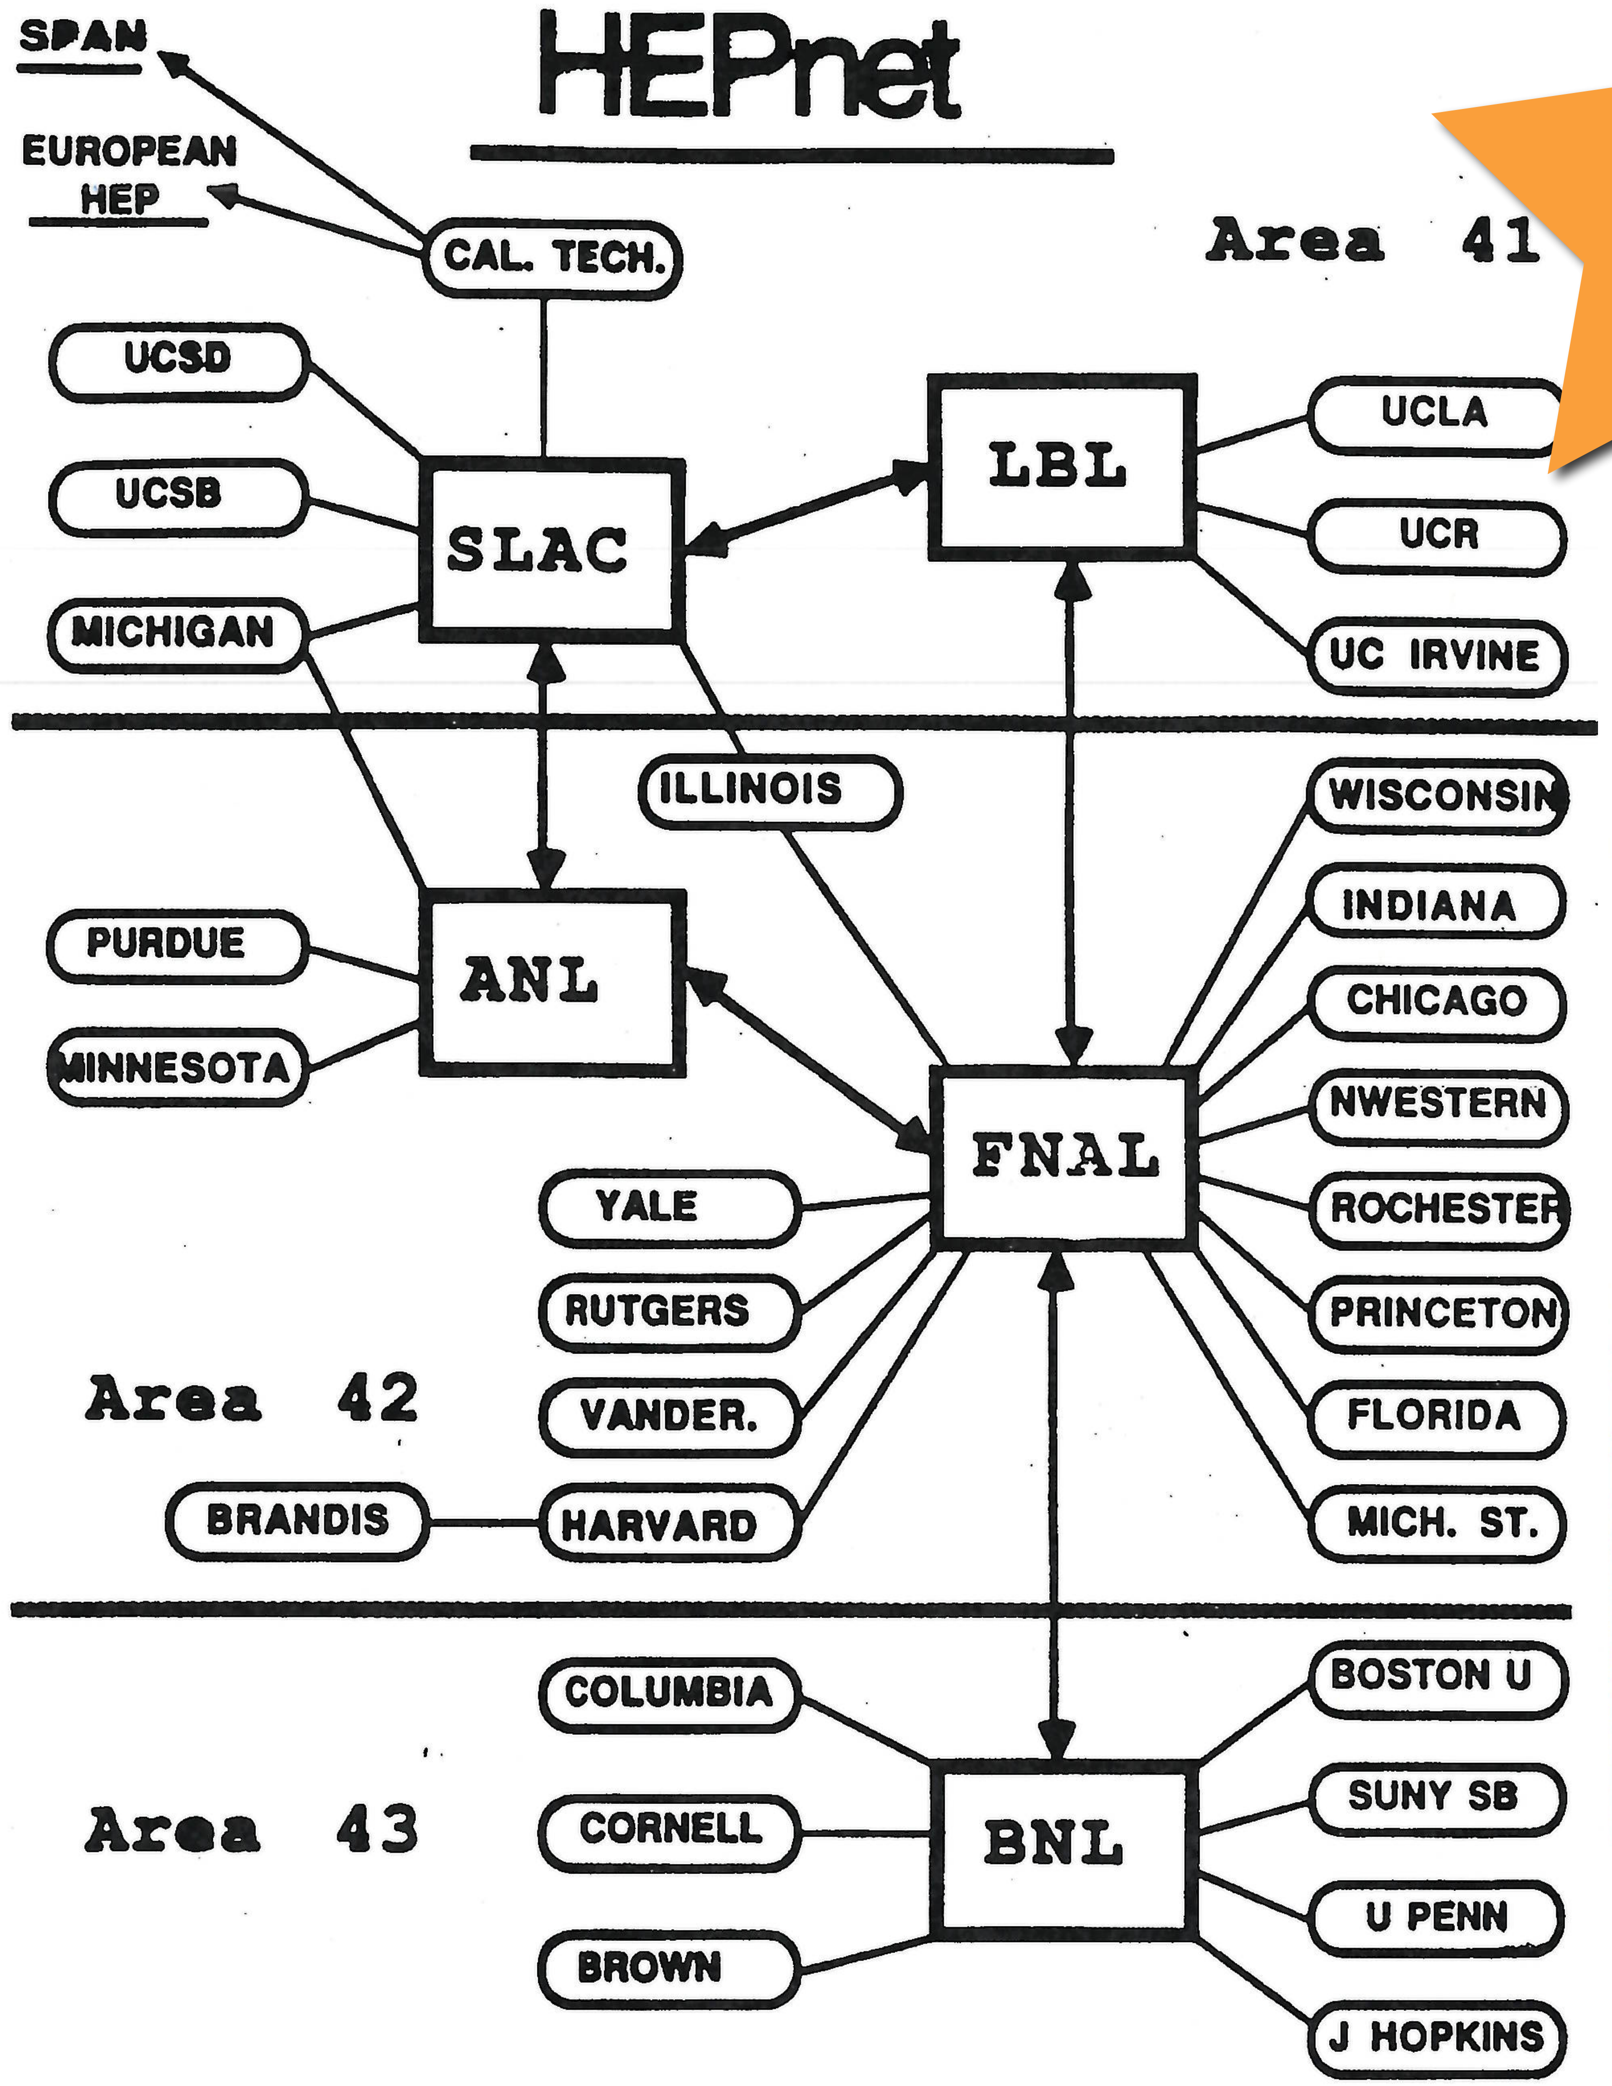

1985

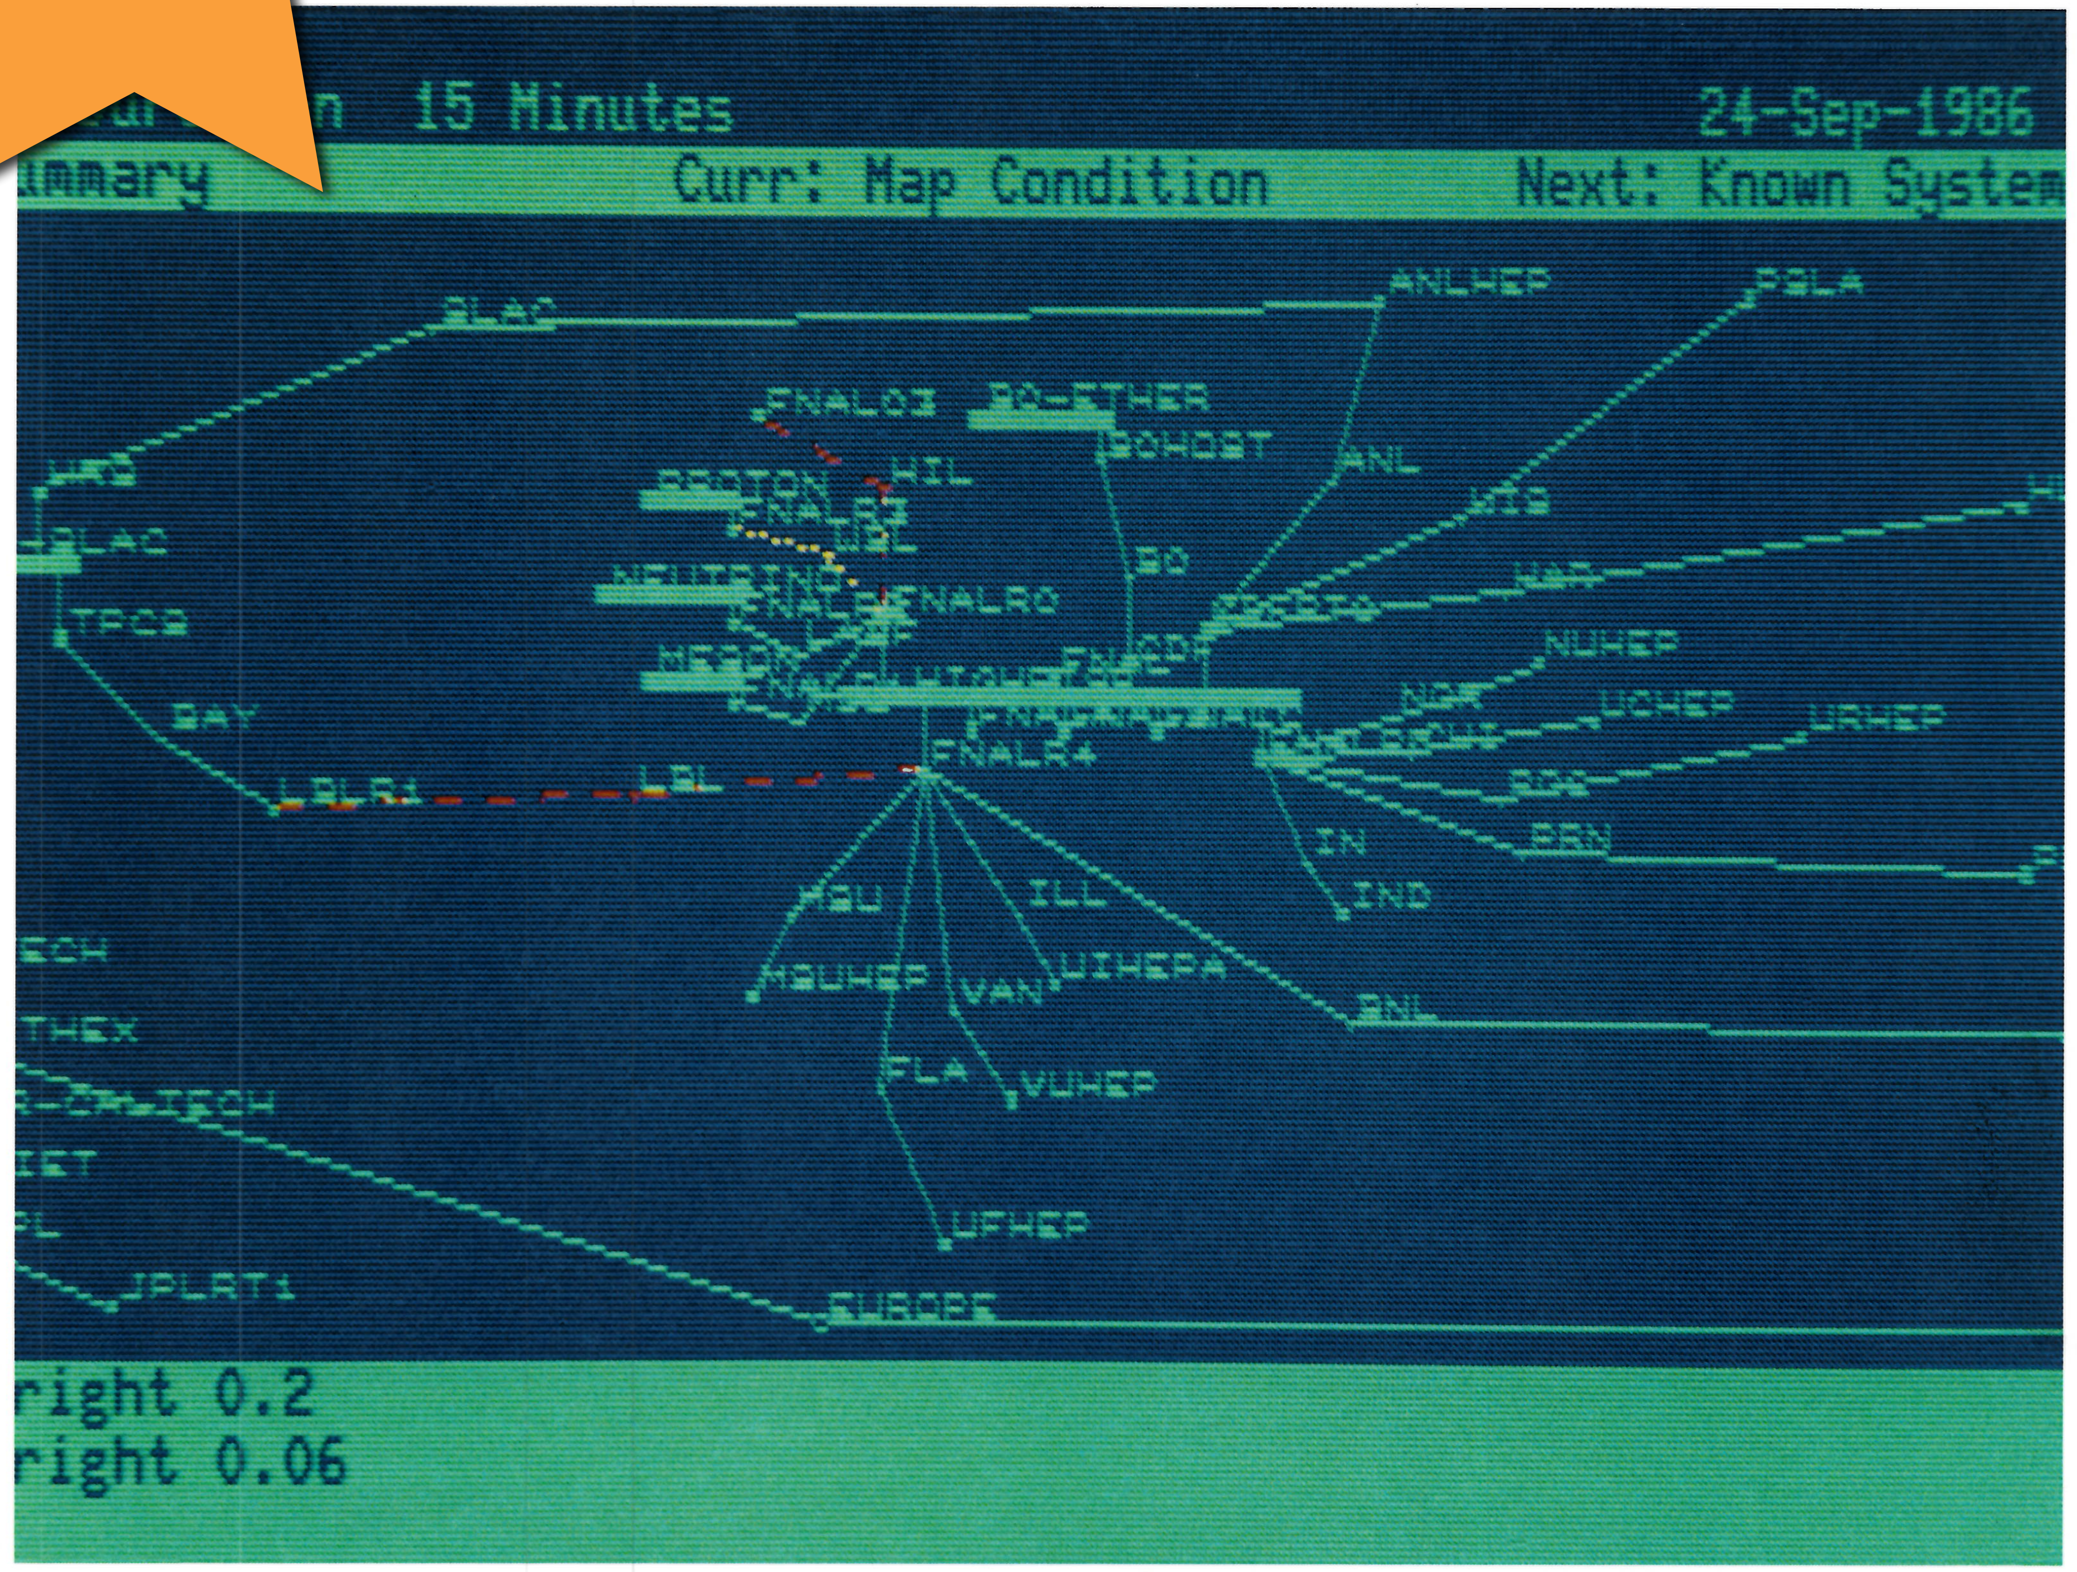

- The world first global network ➜ Fermilab in the center

Fermilab

ACP

Particle Physics is embarrassingly parallel!

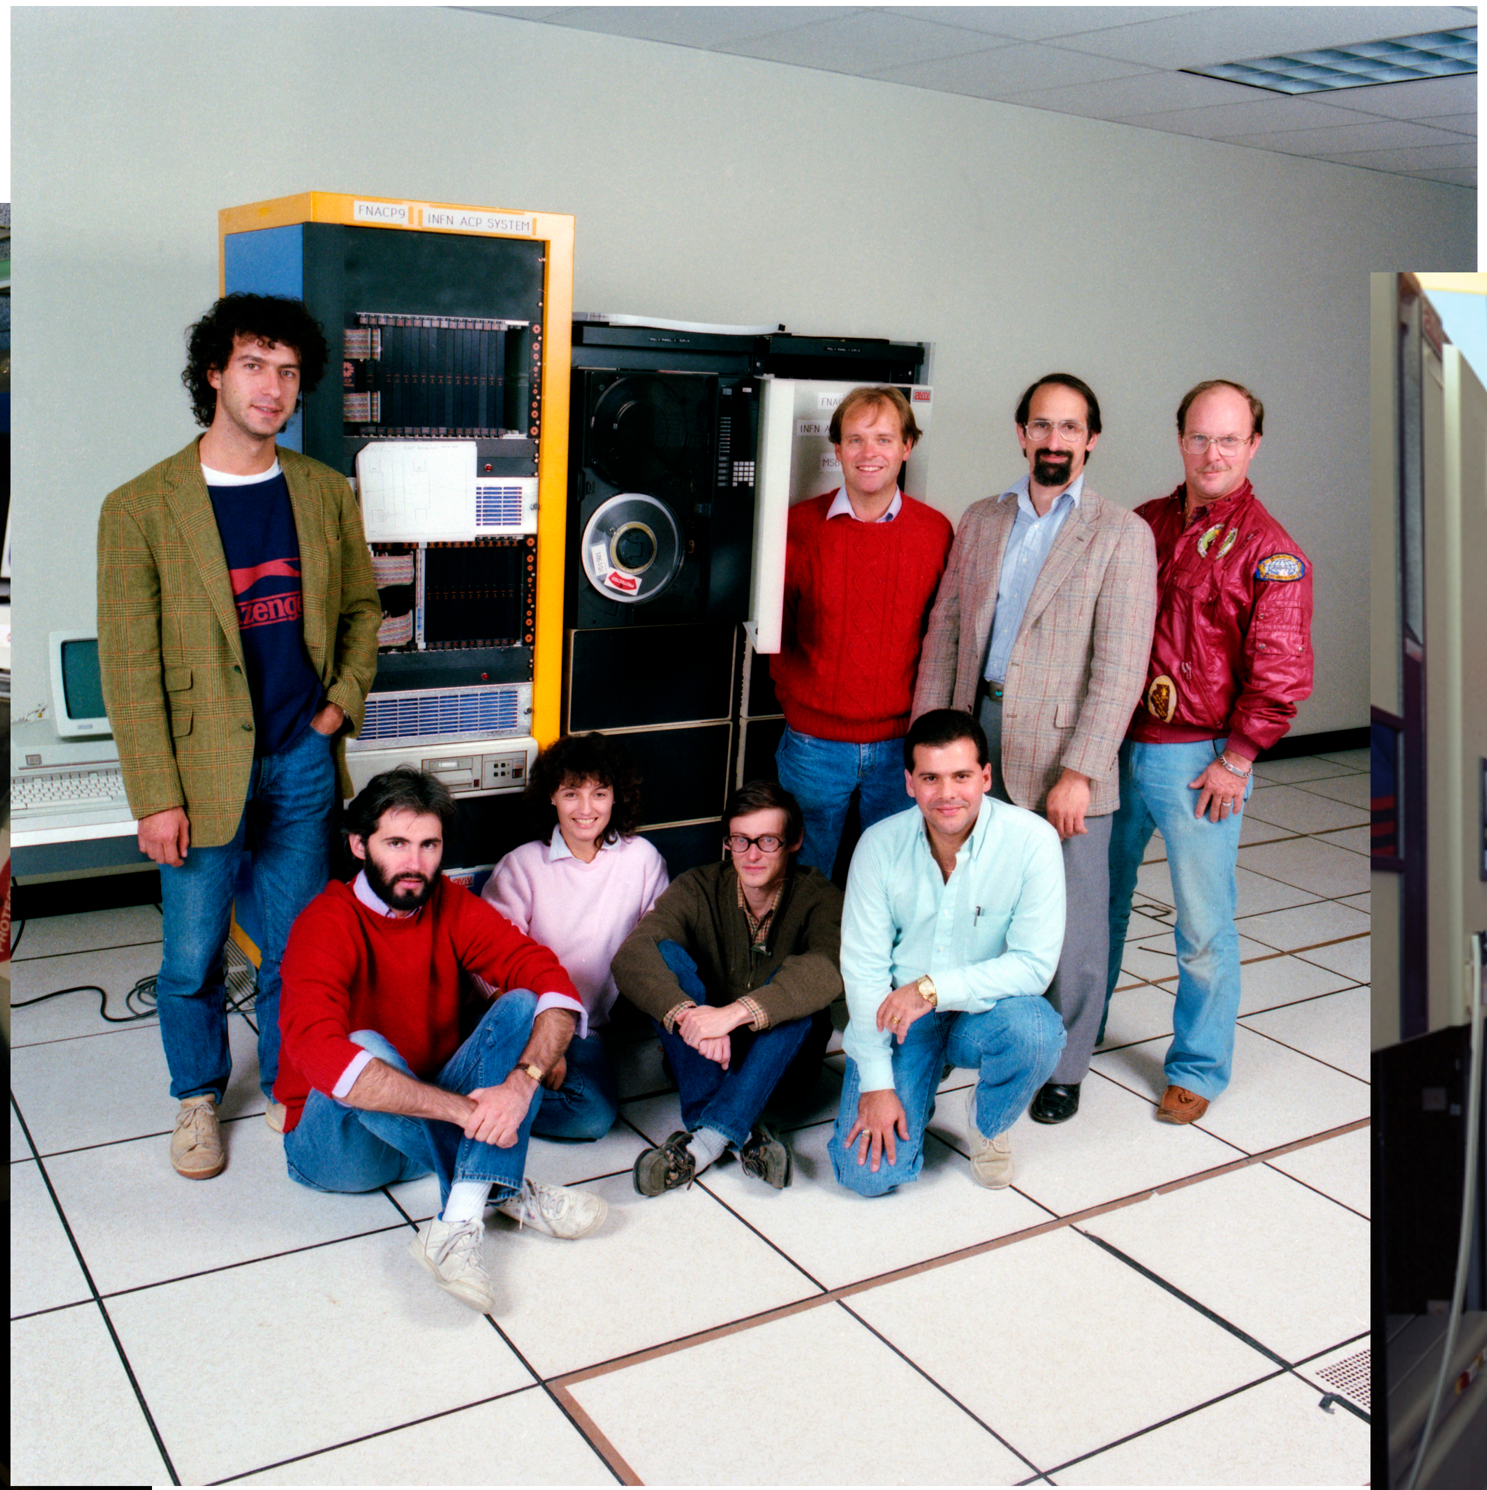

1986

- Starting in 1986, Fermilab built its own boards for parallel processing
  - First for DAQ, then later for trigger (CDF L3 in Run 0) and Lattice QCD (ACPMAPS)

Feynman Computer Center and new mainframe: Amdahl

1988

- first Amdahl in new FCC

Fermilab

Tape robots

1990

- >100,000 8mm tapes in central vault
- 2000 mounts/day average (about 1/2 robotic)
- 1,500 tape drives on site

1994

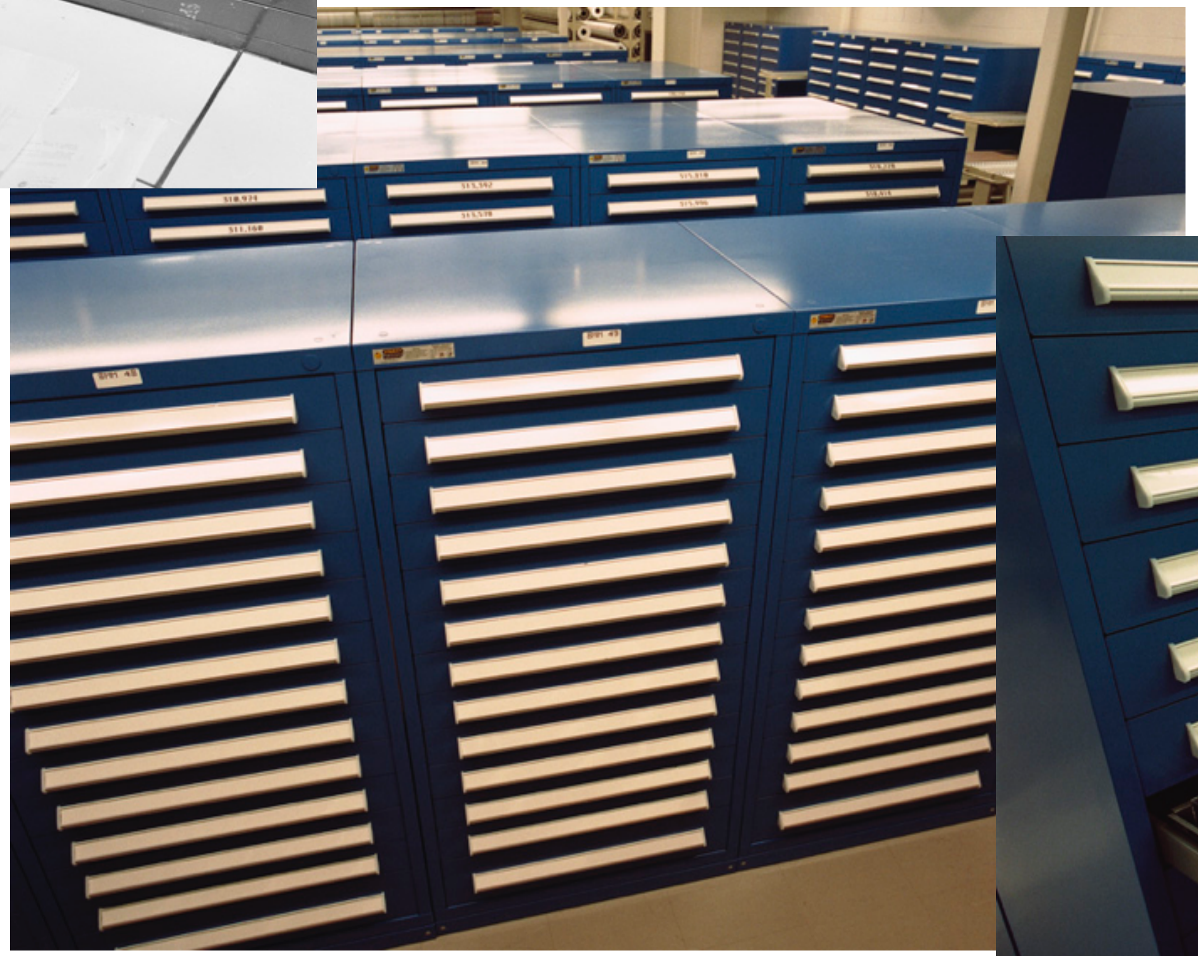

- The switch from manual mounts to tape robots

Common DAQ: Dart

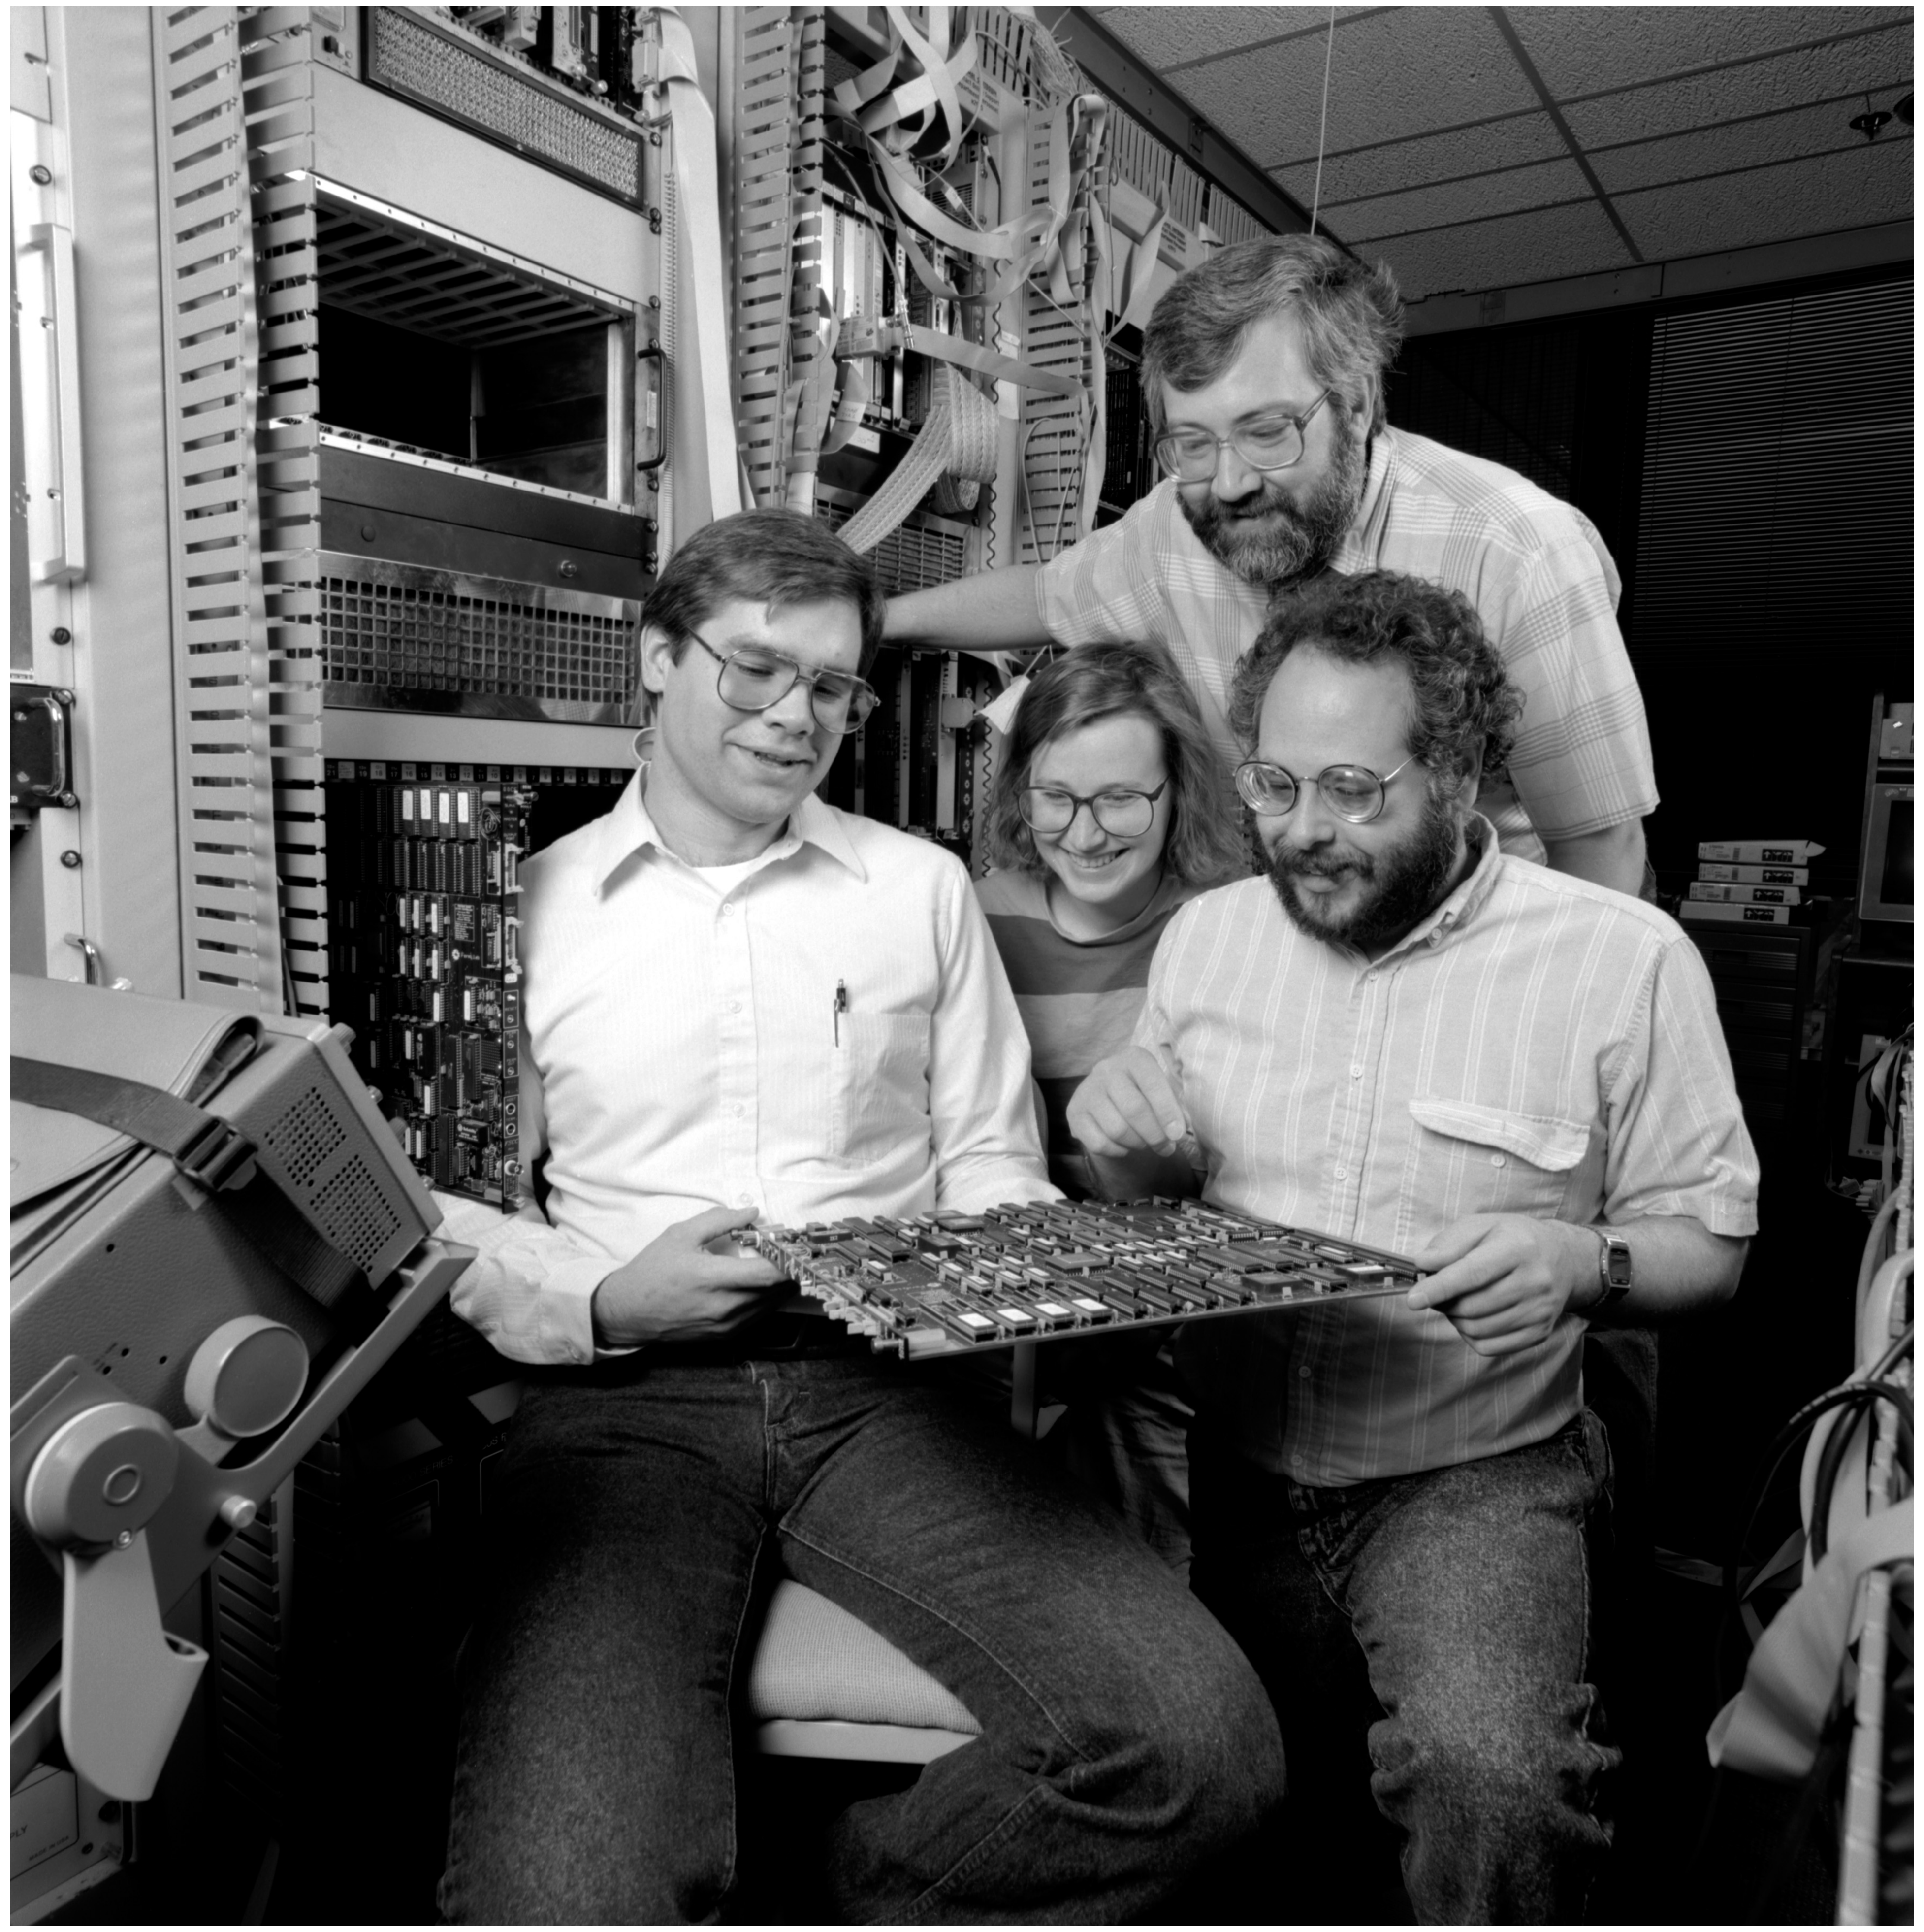

- Dart
  - Innovation: Common DAQ system

The farms - LINUX

1997

1998

- 1997: PC Farm pilot project
- 1998: First PC Farm

Fermilab

Mystery in Computing - Robots developing self-awareness???

- Robot broke glass door in the night
  - Pretended nothing has happened and continued mounting tapes
- At first no explanation:
  - Uprise of the robots?
- In the end:
  - Outgassing in the tempered glass caused the door to shatter

The software evolution(s) and revolution(s)

- 1996:
  - Tevatron Run 2 upgrade project: Fortran ➜ C++
- 1998:
  - ROOT selected as the Fermilab supported analysis tool
- 2006:
  - Fermilab team replaces CMS software framework with CMSSW
- 2010:
  - *art*, the intensity frontier software framework, started as a fork of CMSSW

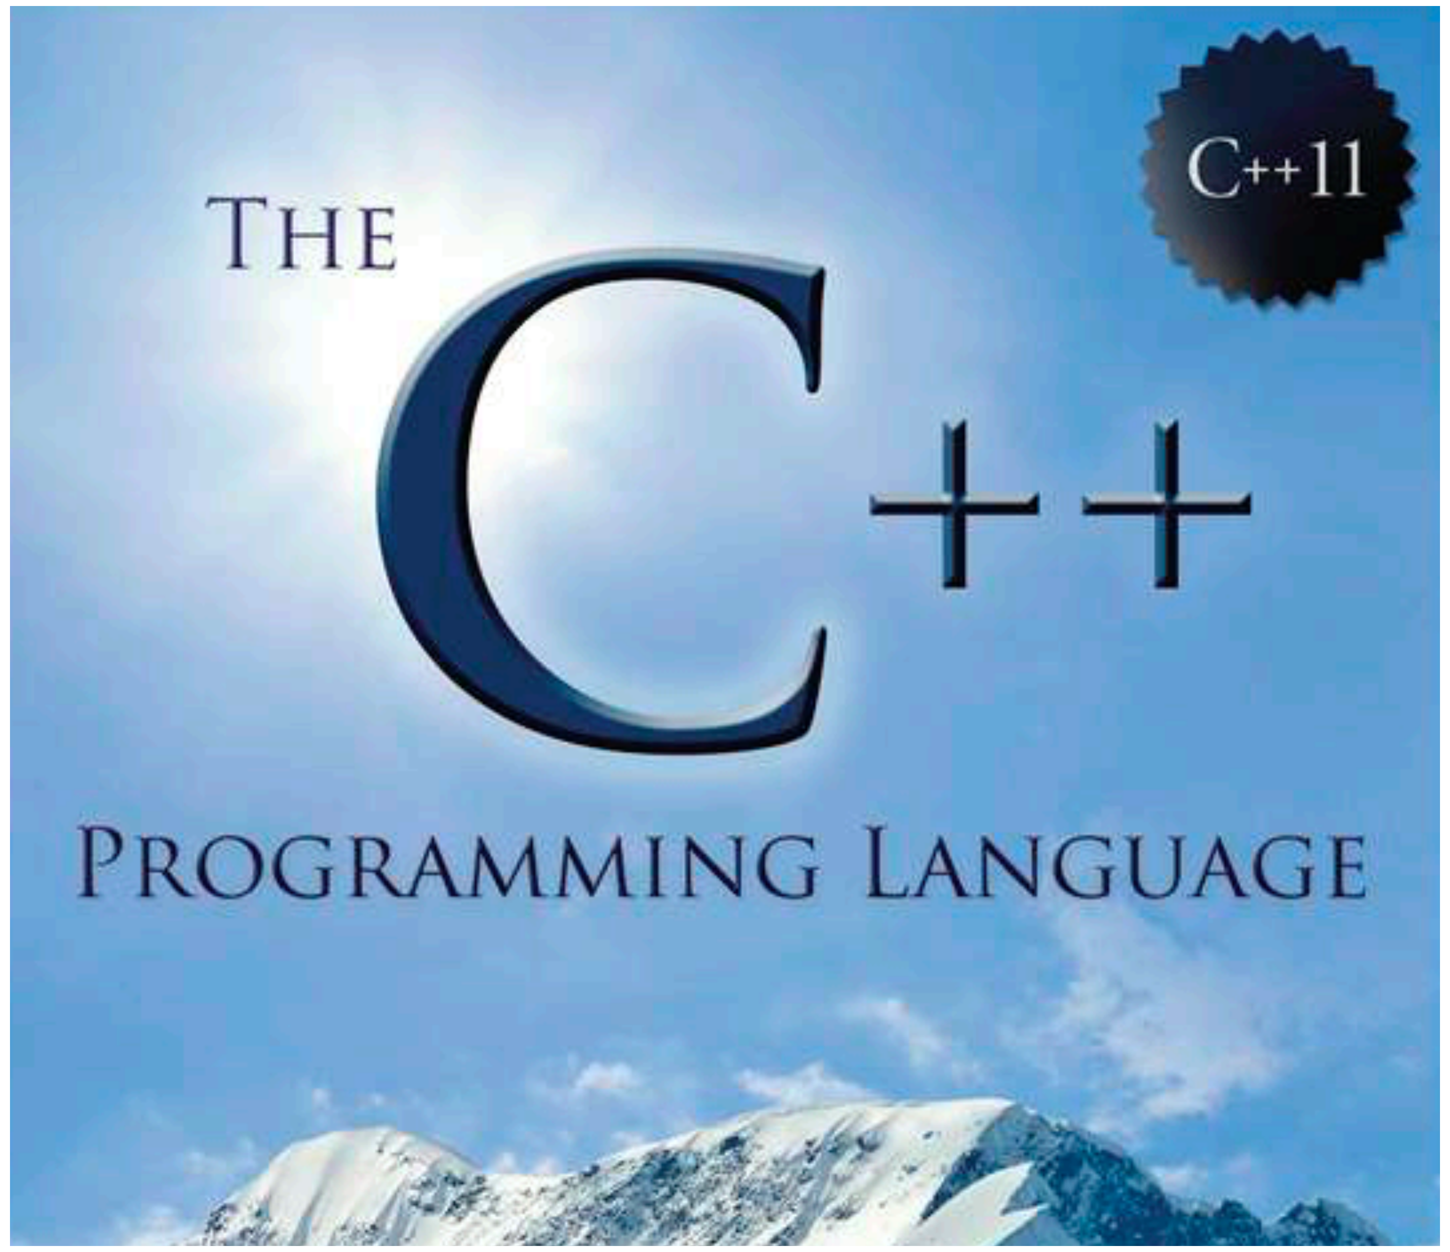

ROOT
Data Analysis Framework

Fermilab

Sloan Digital Sky Survey

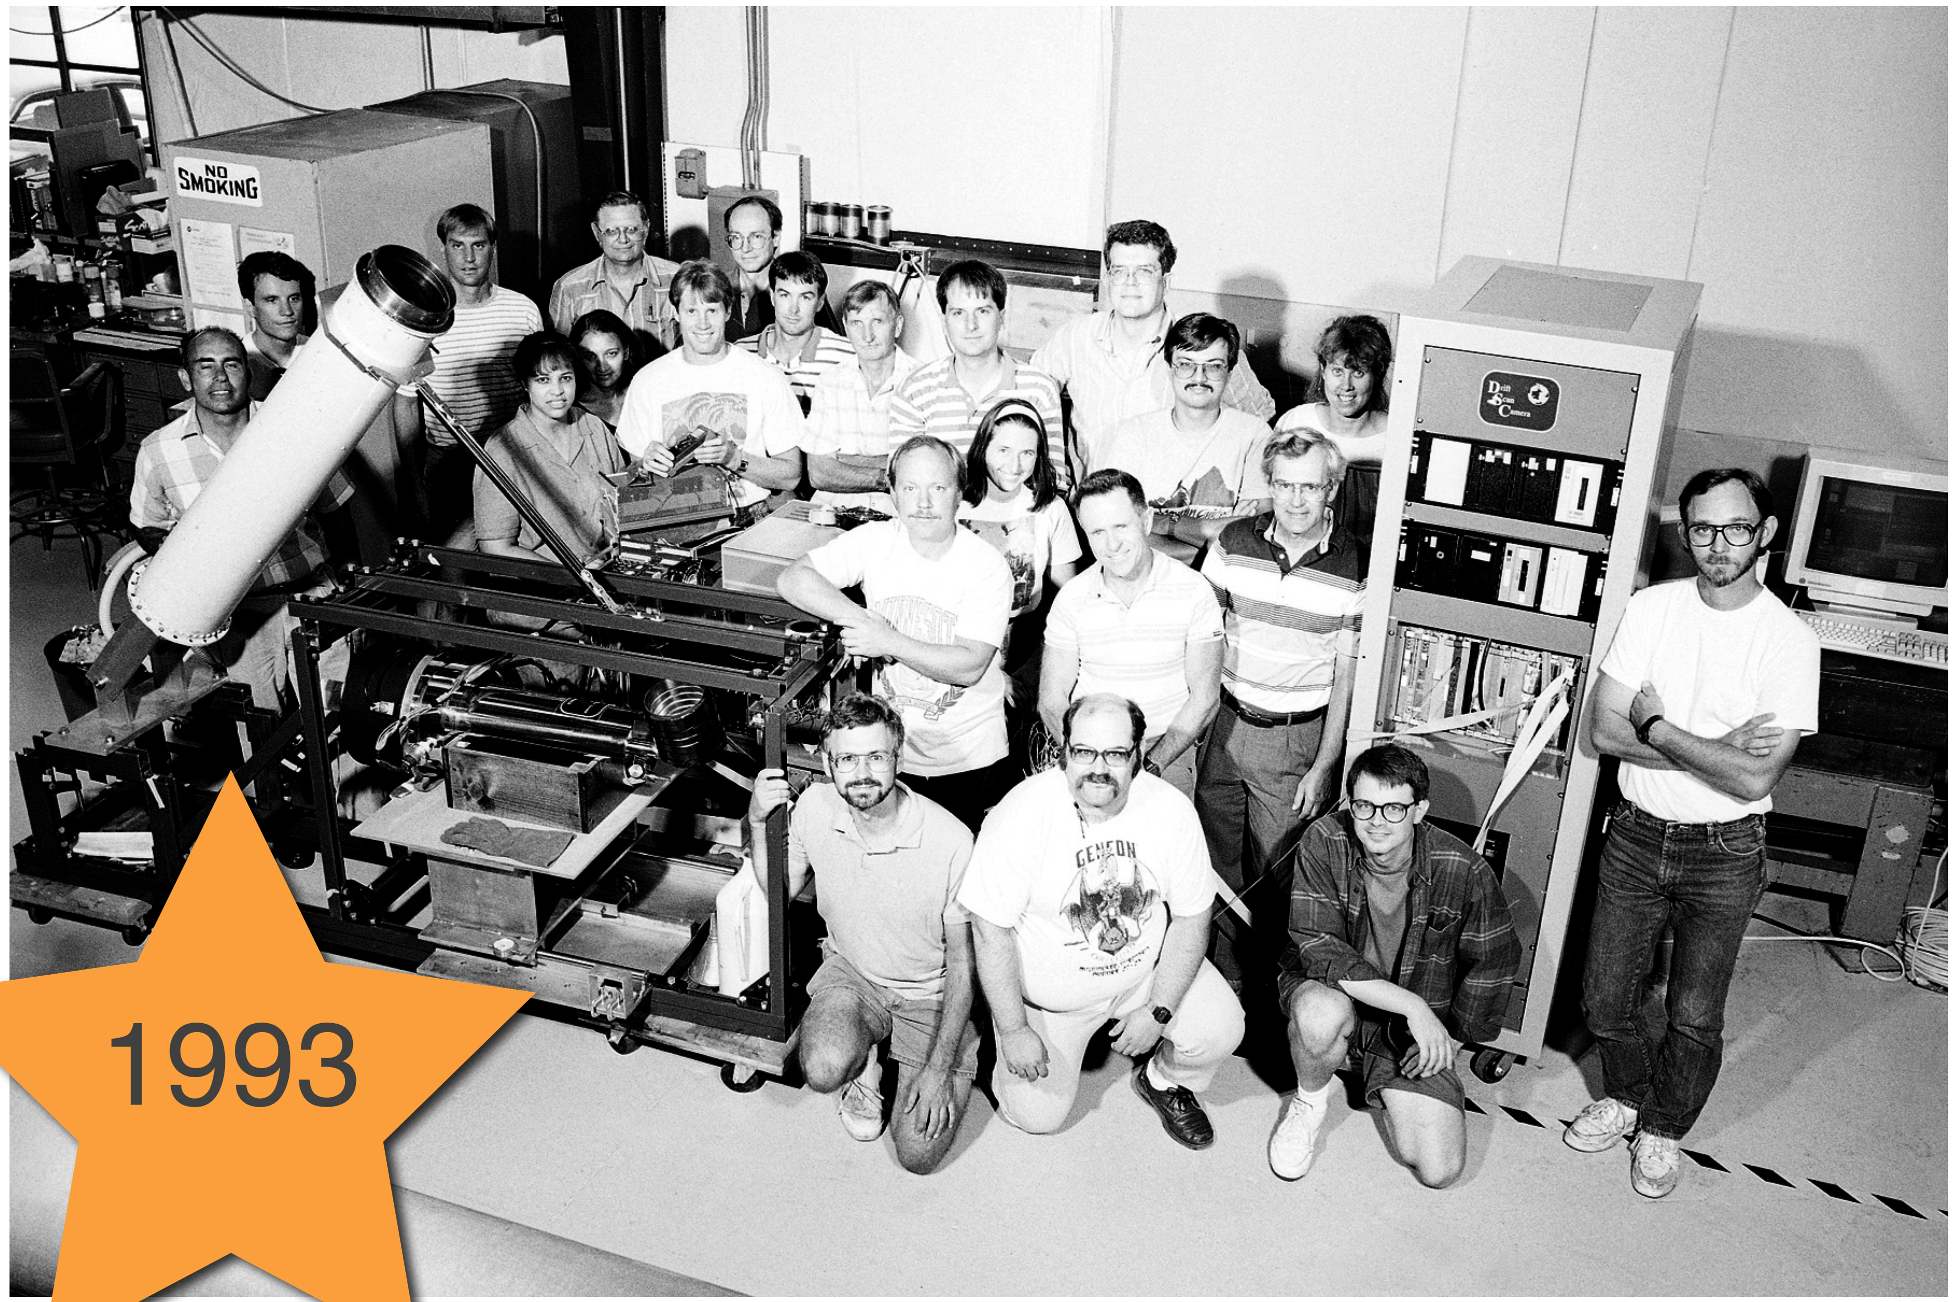

- Contributed HEP expertise to Astro-Particle Physics
  - DAQ development
  - Software development, standardized tools
  - Data handling and processing

SciDAC

2001

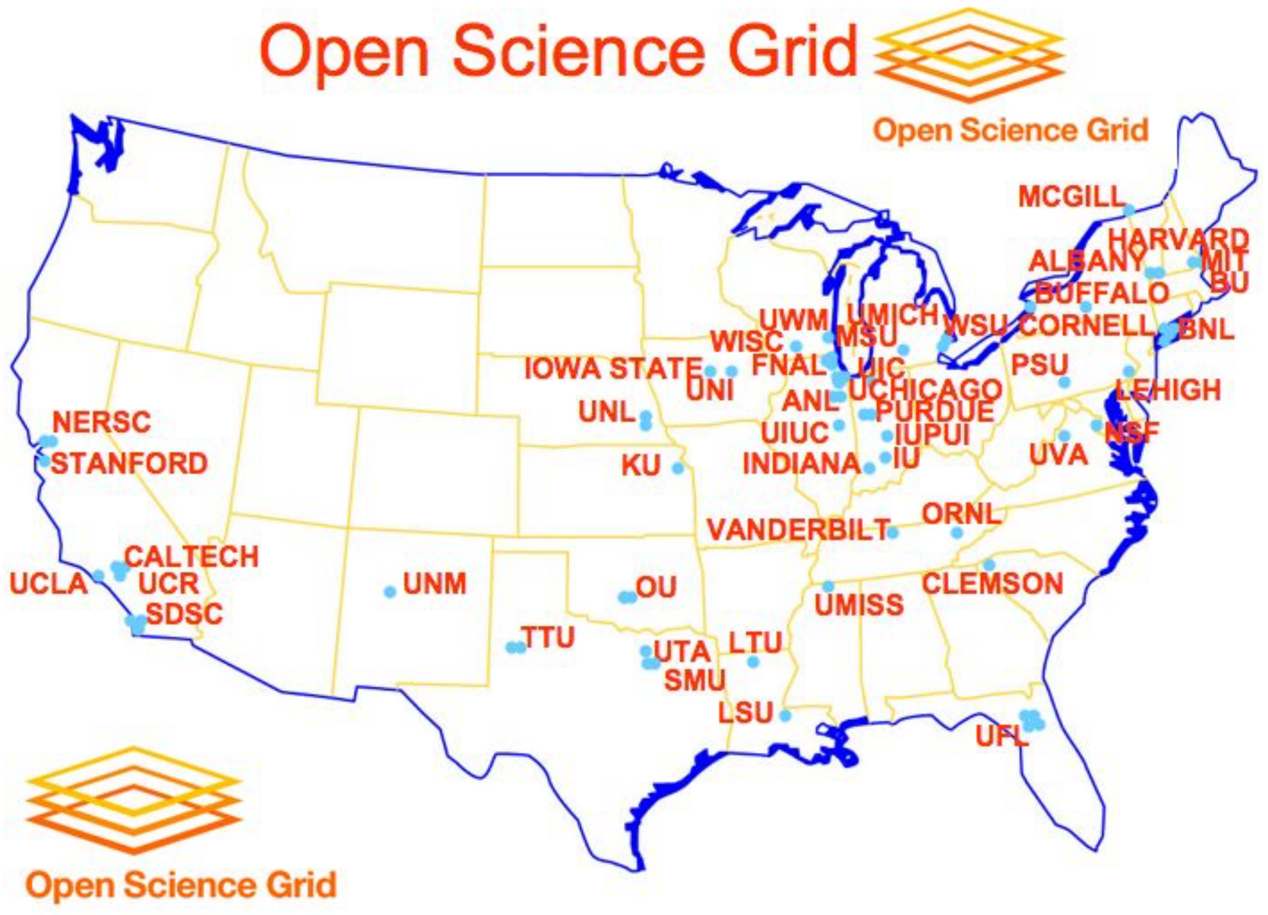

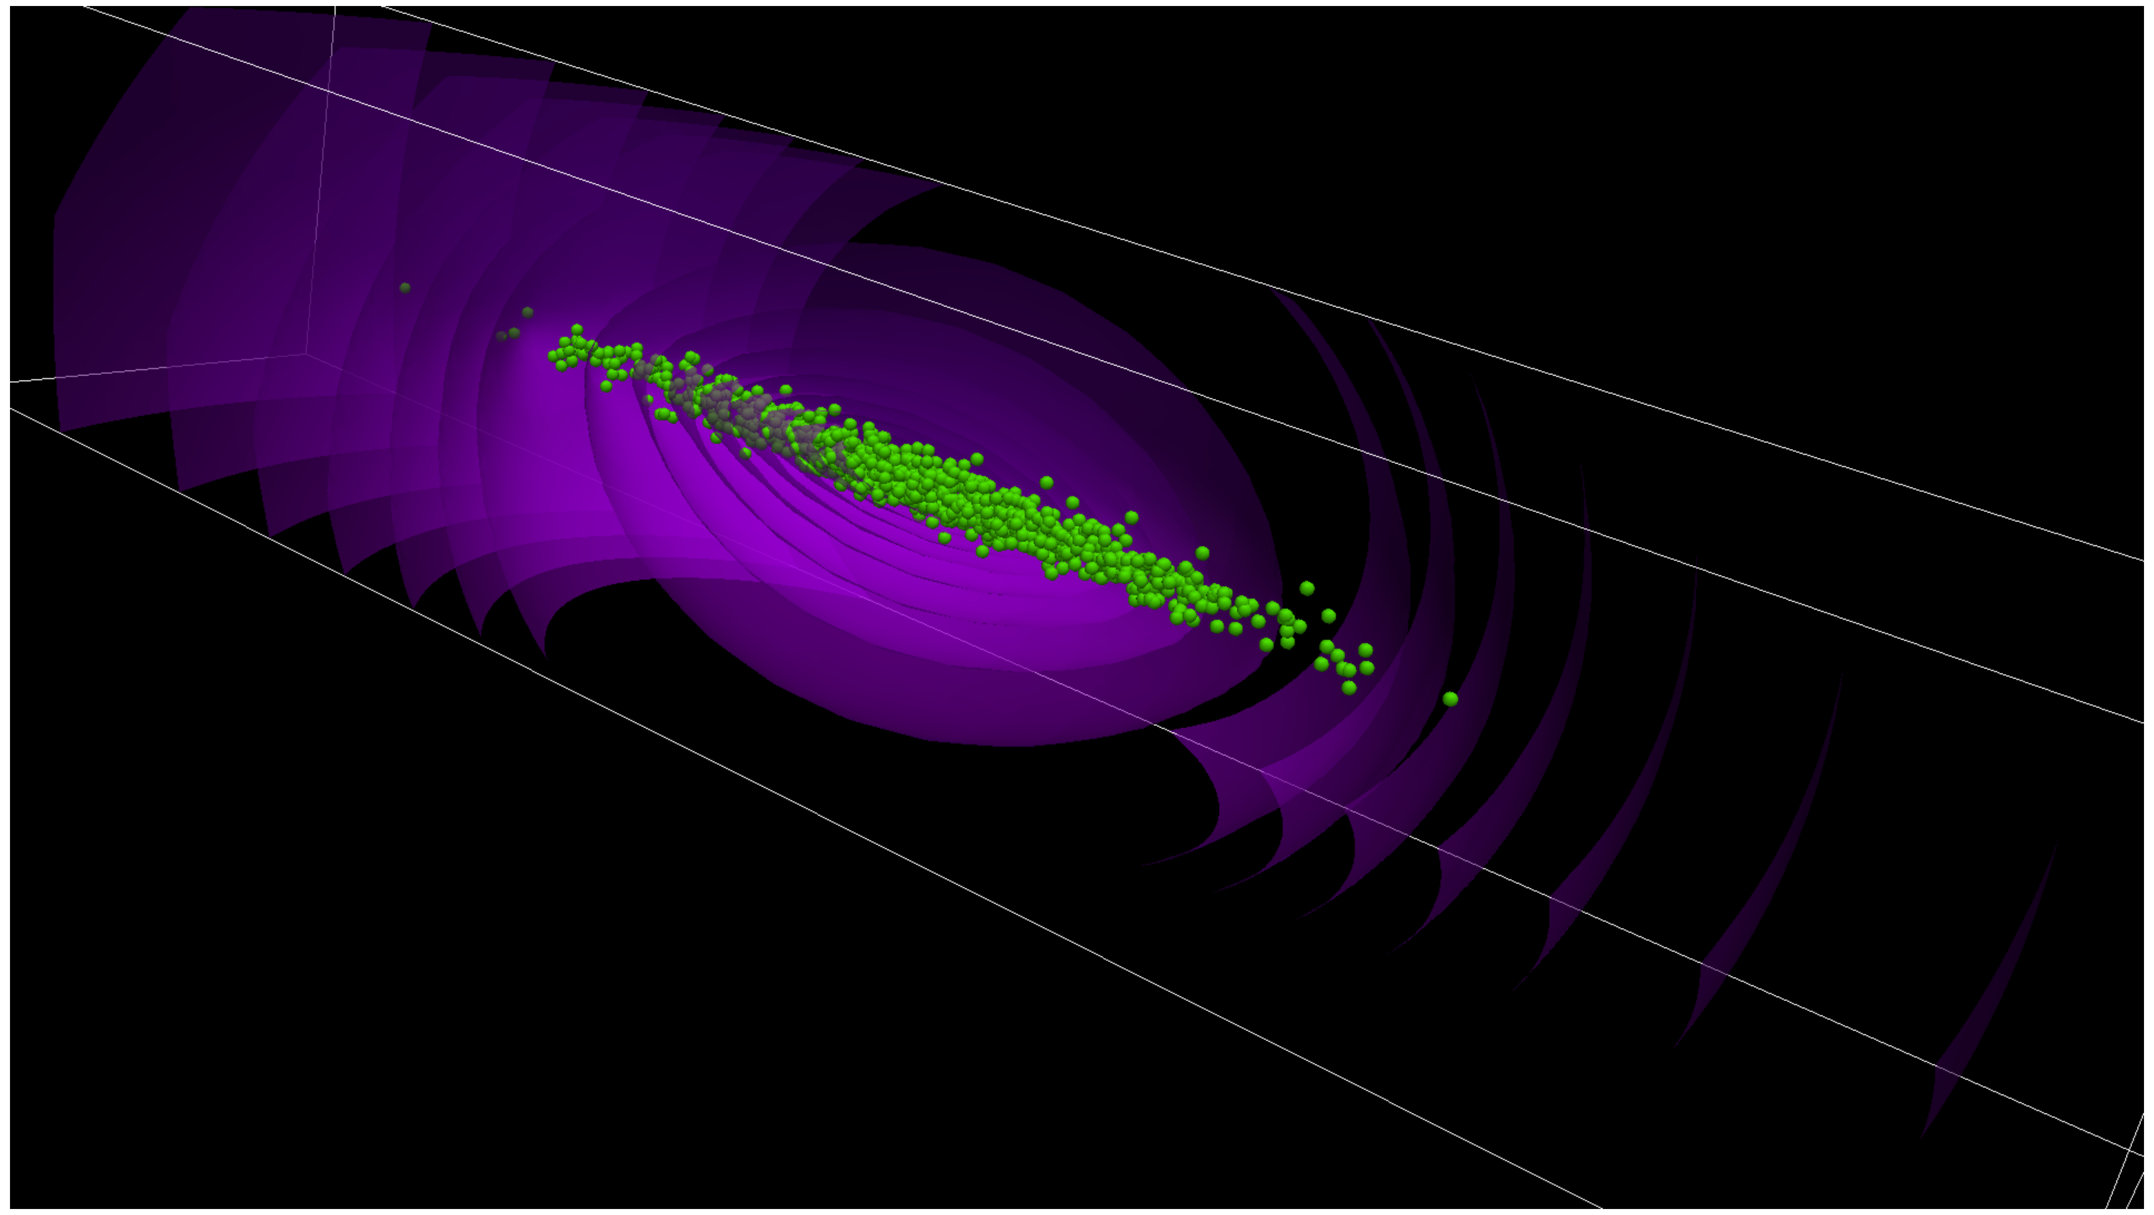

Lattice Gauge Computing

Open Science Grid

Accelerator Science and Simulation

- Fermilab receives DOE SciDAC awards for work in Lattice Gauge Computing, Open Science Grid and Accelerator Science and Simulation

LHC computing

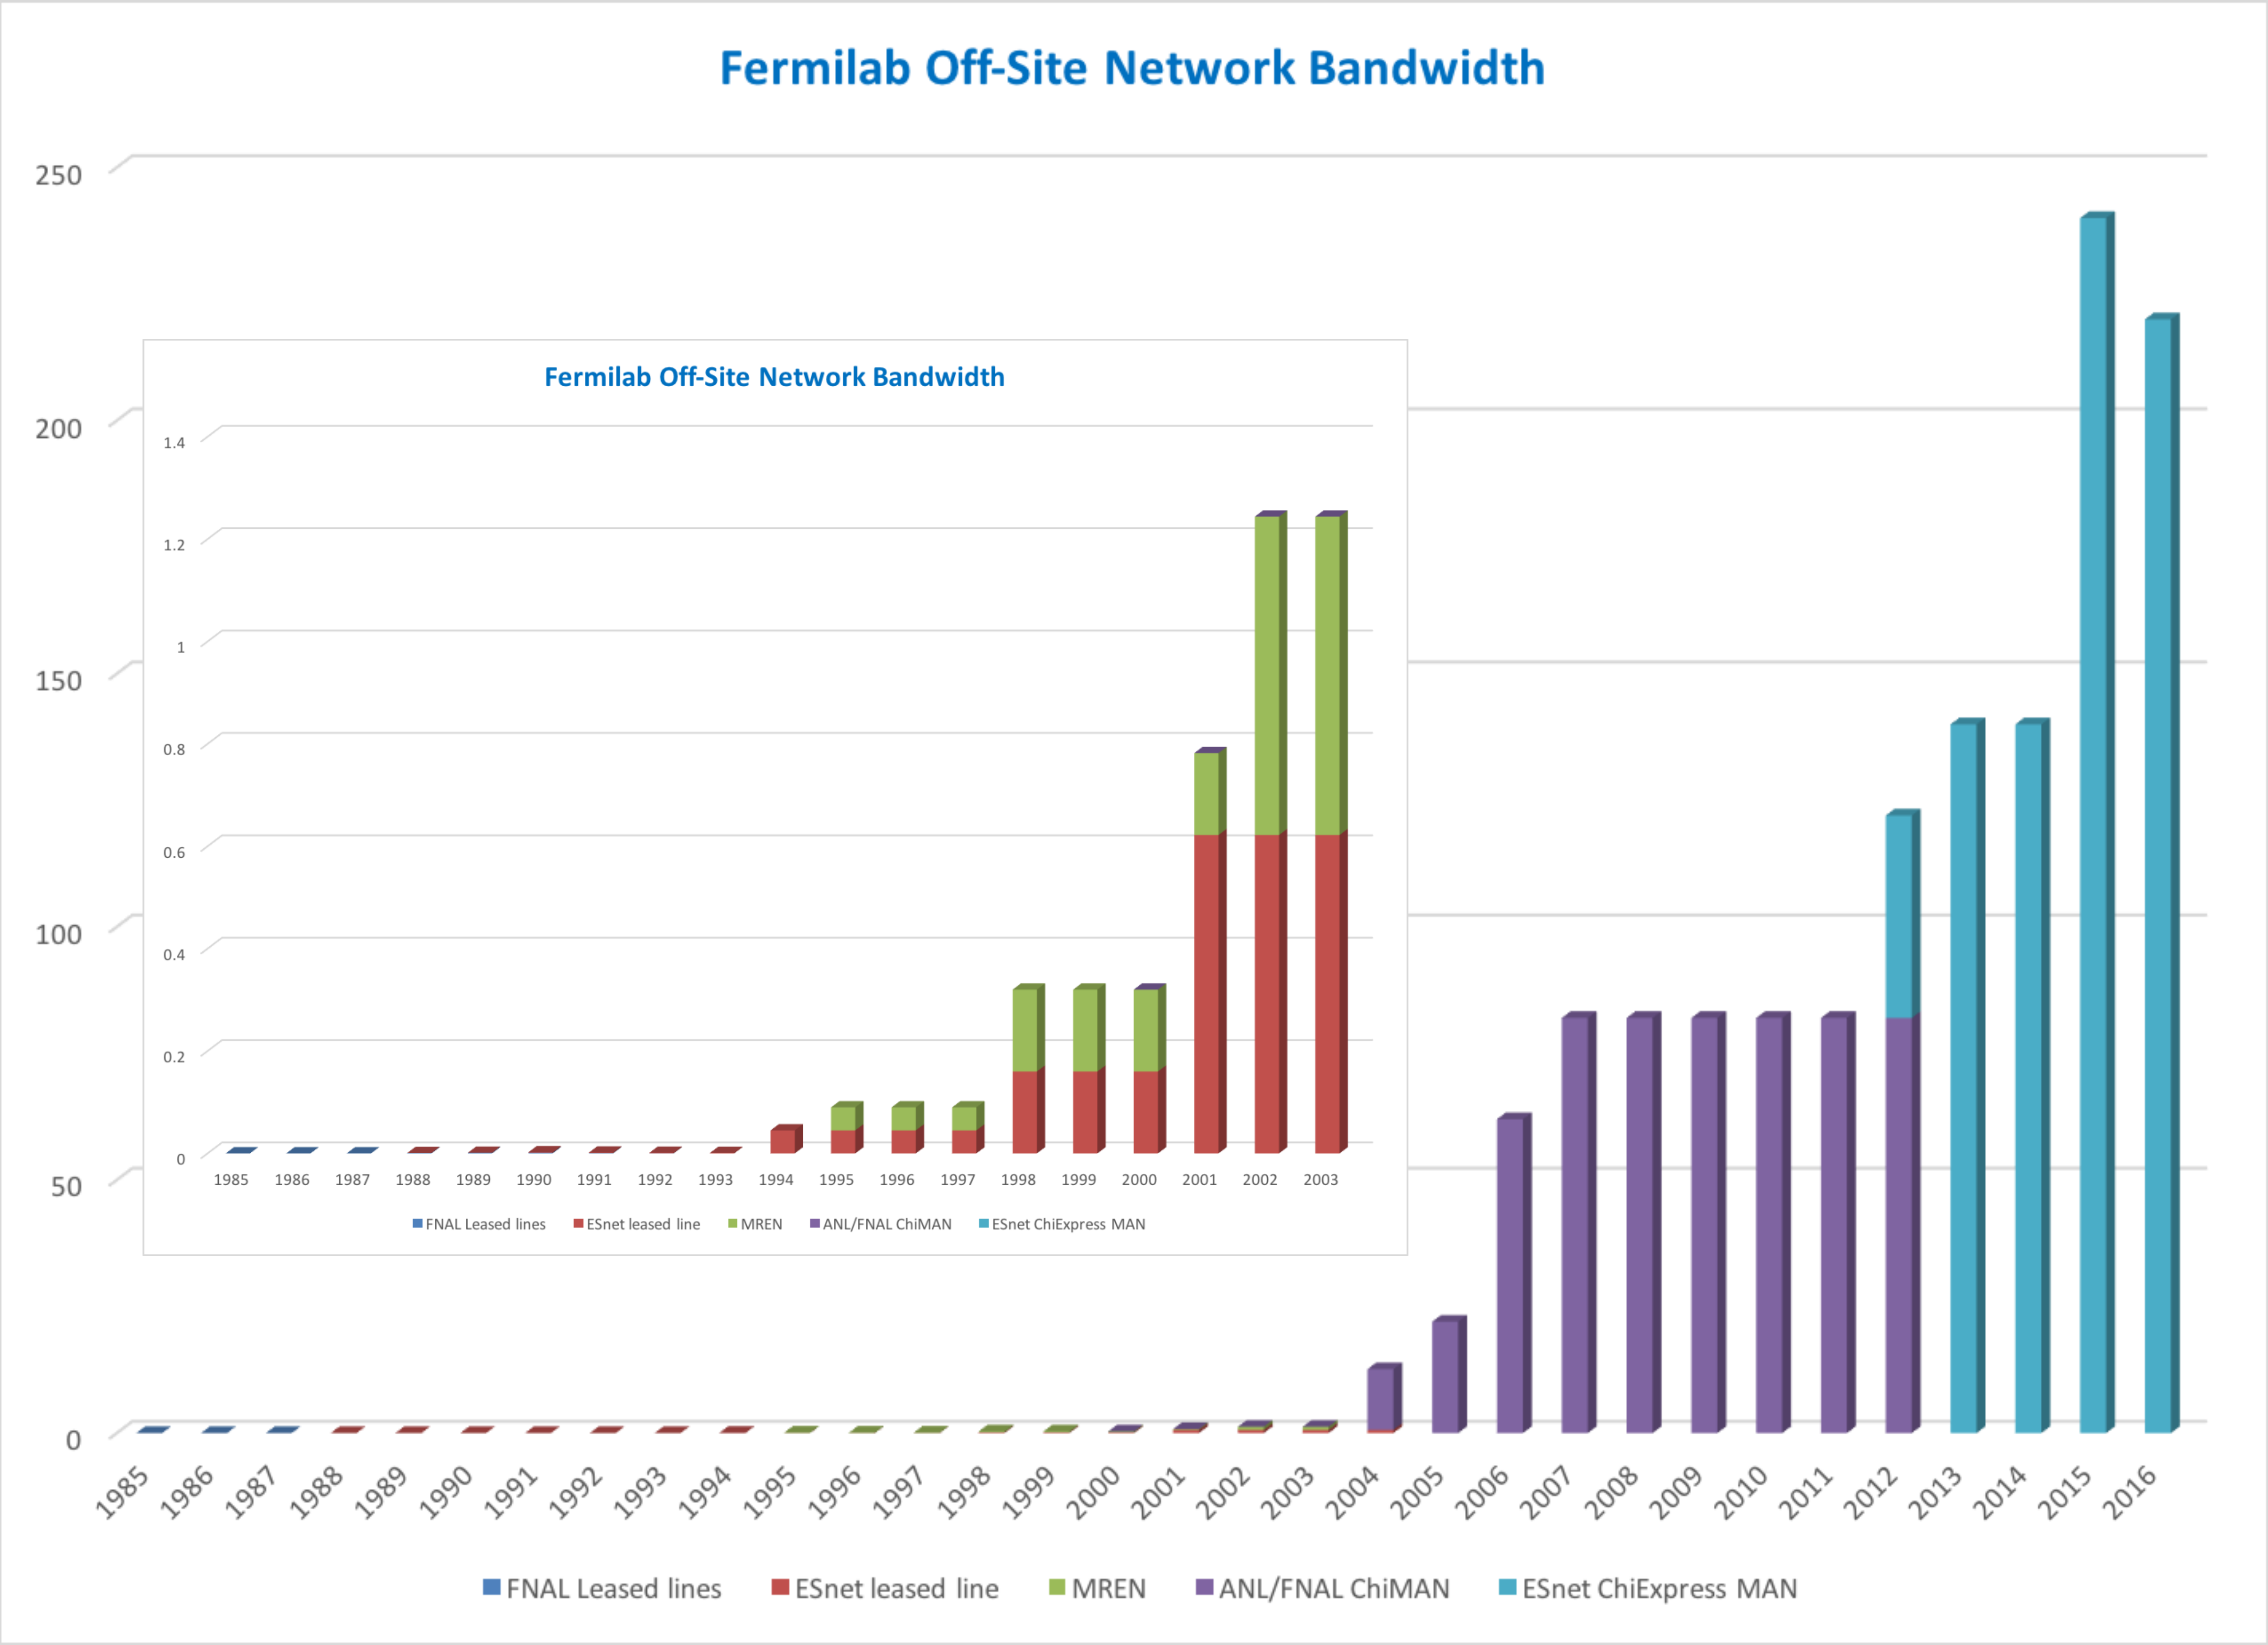

2001

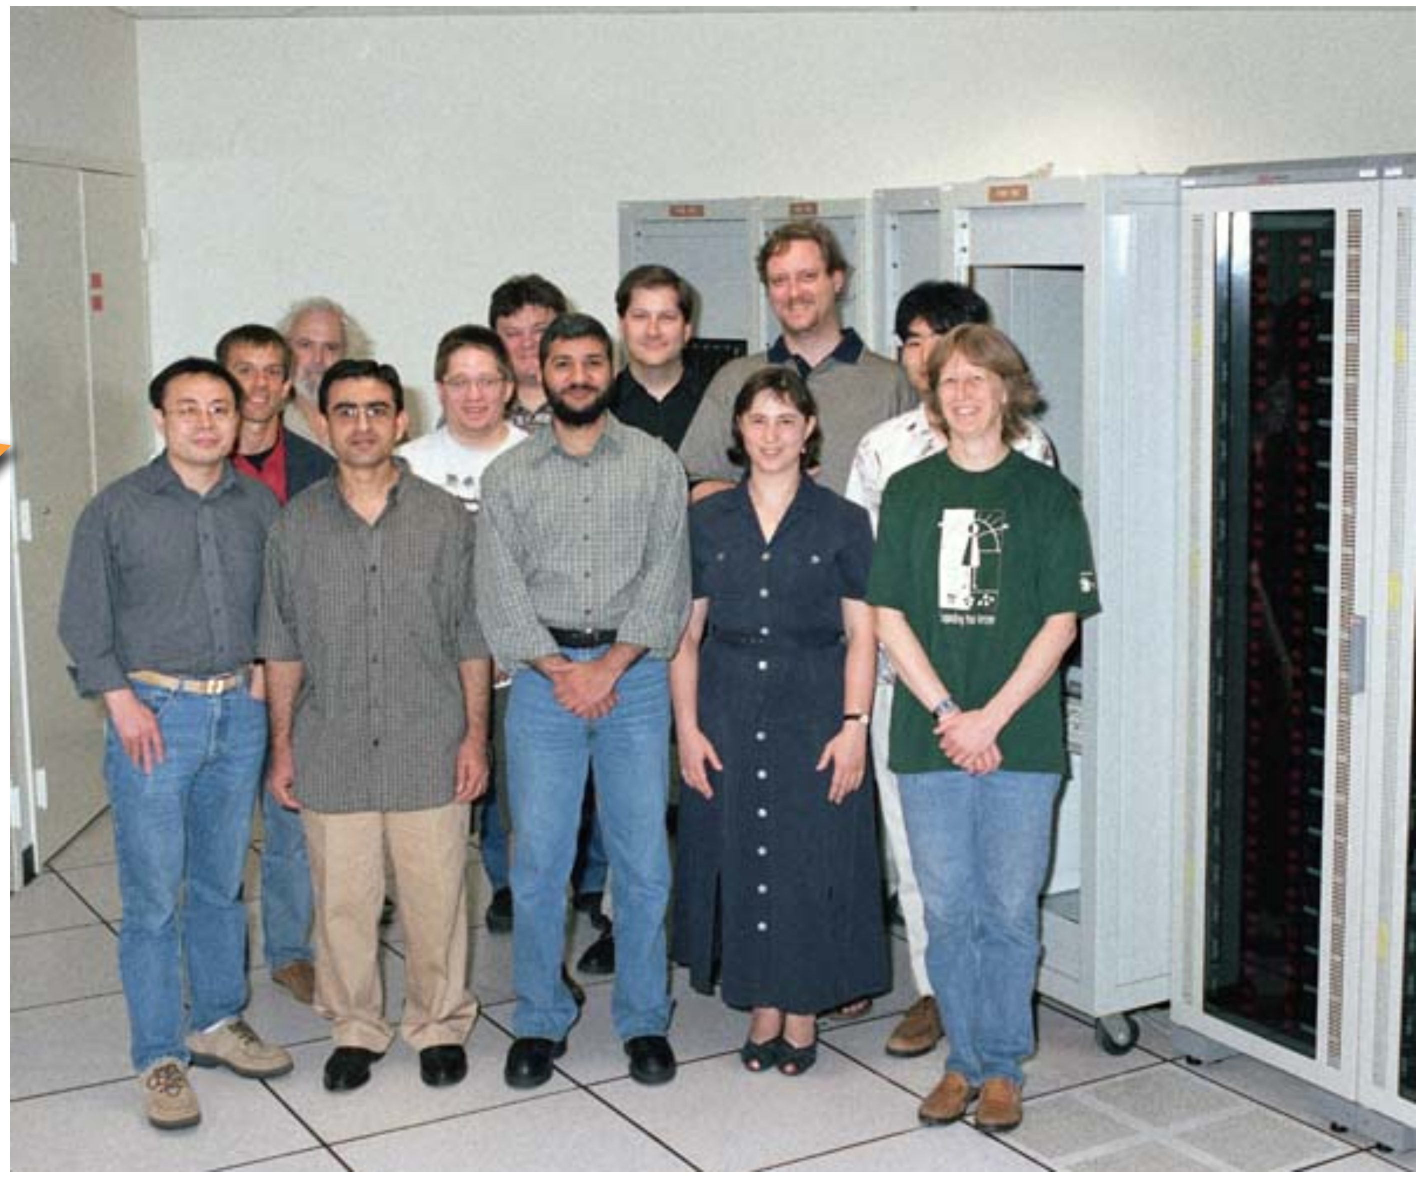

2005

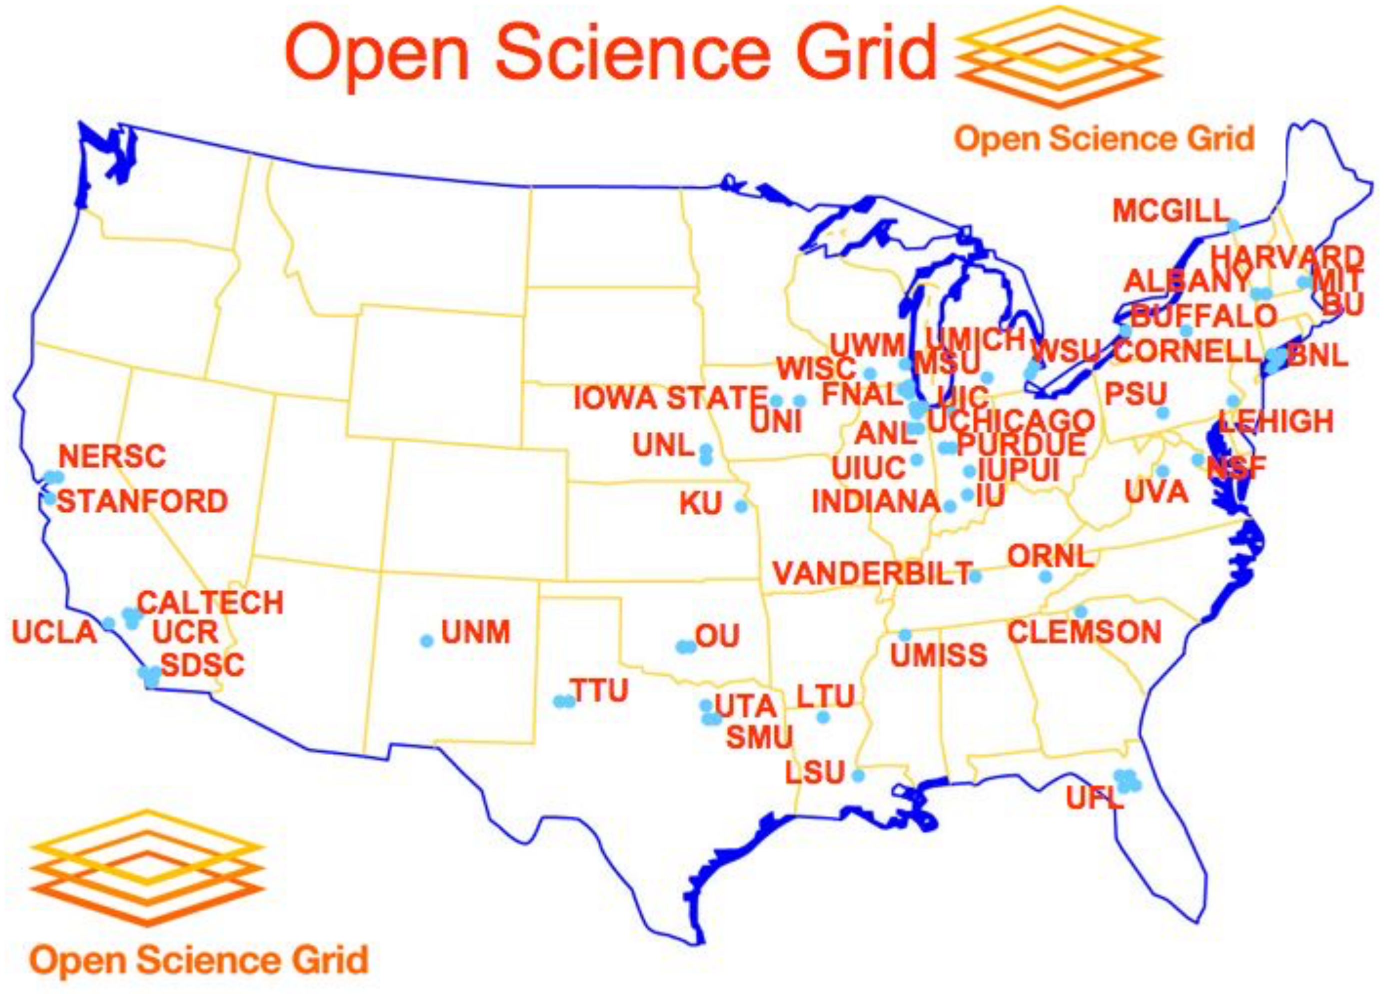

- USCMS Tier-1 site at FNAL established in 2001 and LHC transatlantic networking starting with USLHCNet grant to Caltech
- Open Science Grid established in 2005

The present

CPU

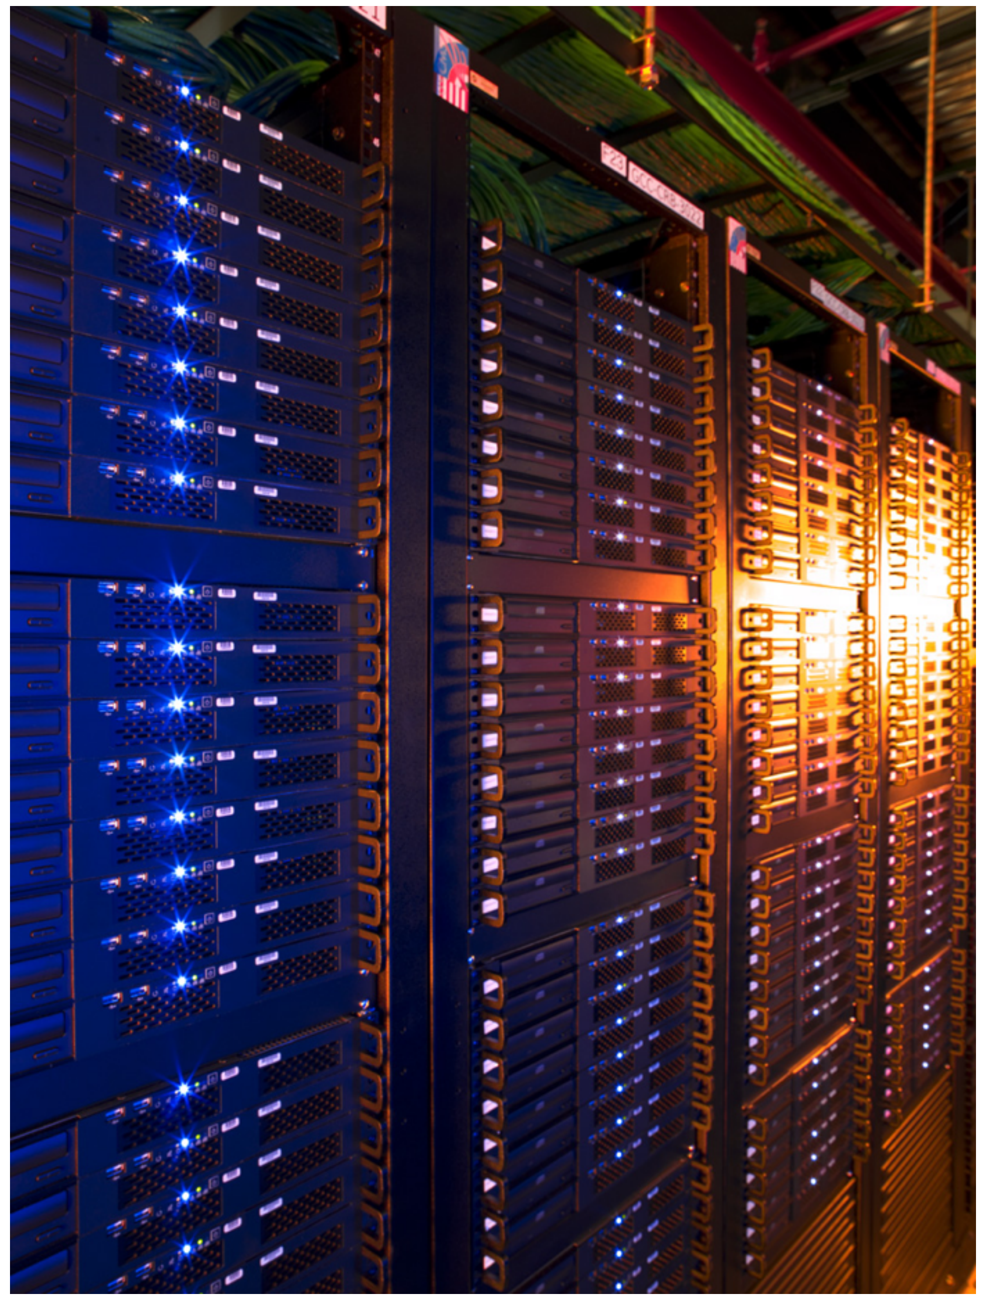

48,000 cores
(plus 20,000 HPC cores)

Tape

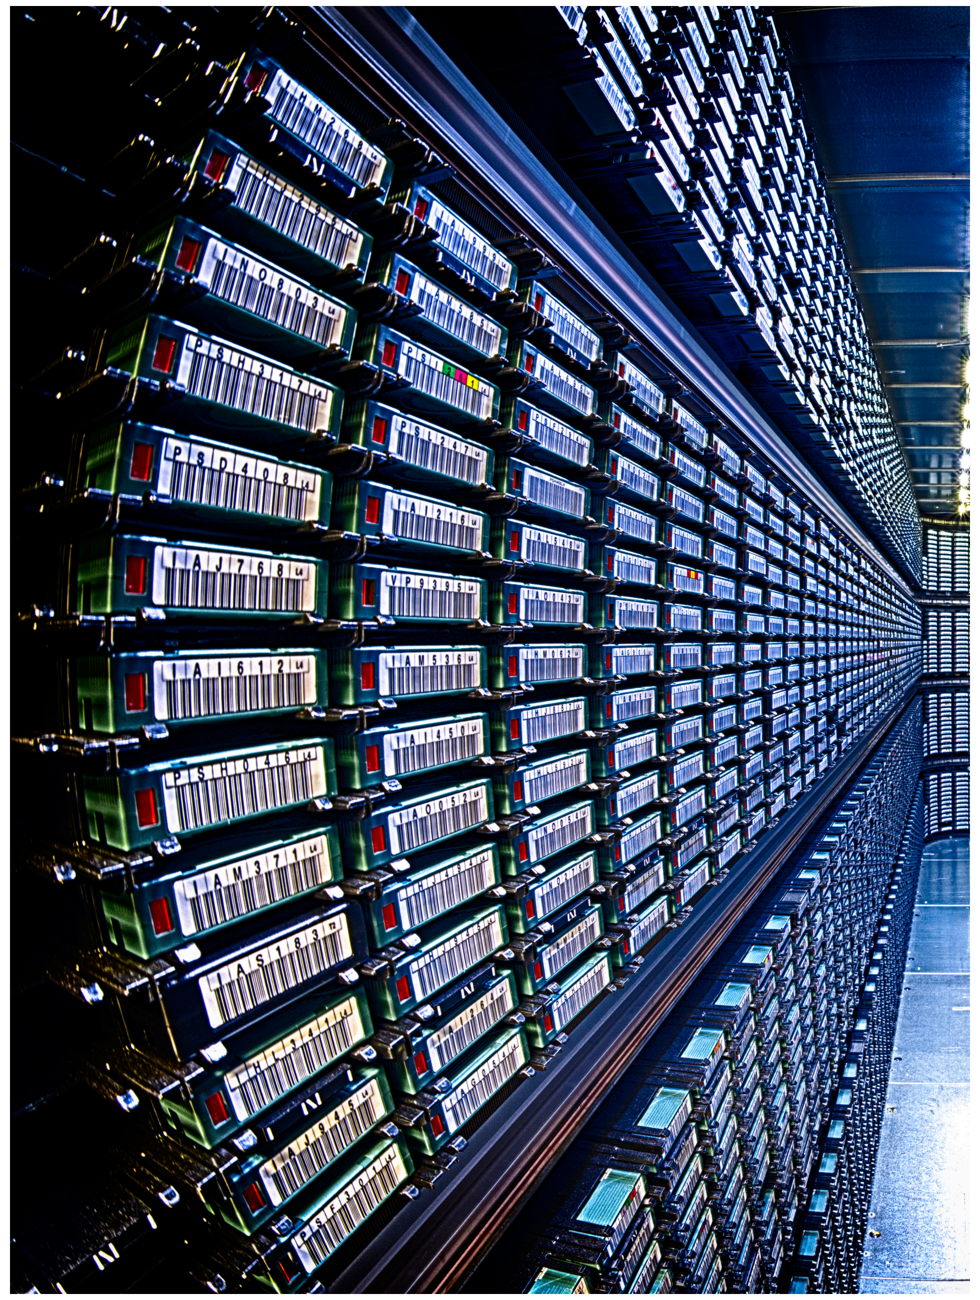

1 Exabyte Capacity
100 PB of tape media, 90 PB used

Disk

35 PB spinning disk

Network

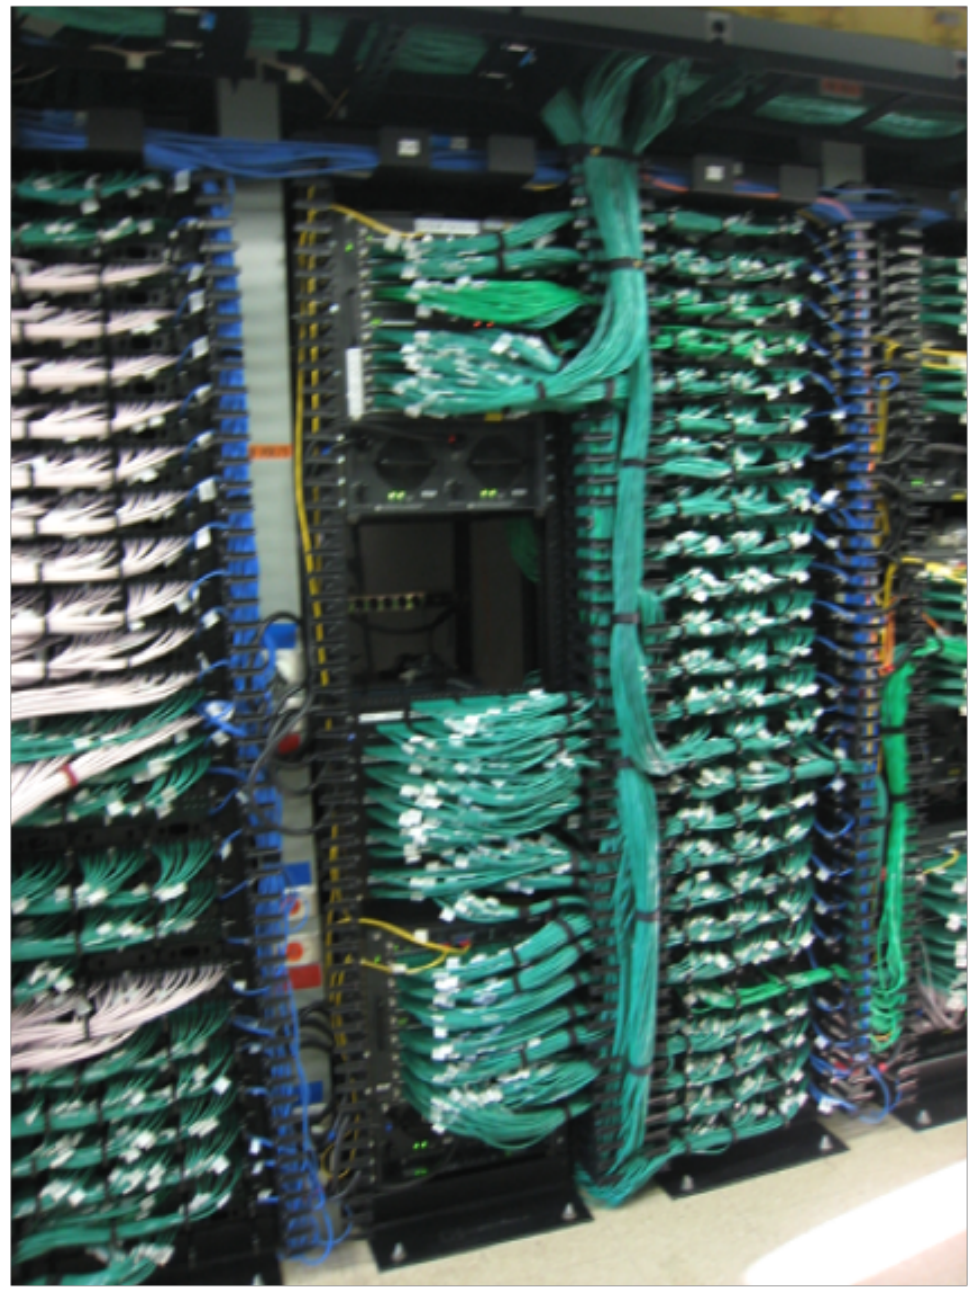

2x100 Gb/s offsite
~35,000 internal network ports

Open Science Grid

WLCG
Worldwide LHC Computing Grid

Common Tools

- Software Frameworks
  - CMSSW
  - art
- Reconstruction Software
  - LArSoft
- Grid
  - glideinWMS
  - FIFE
- Simulation
  - Pythia
  - Geant
  - Synergia

The Future: Massively Parallel Software

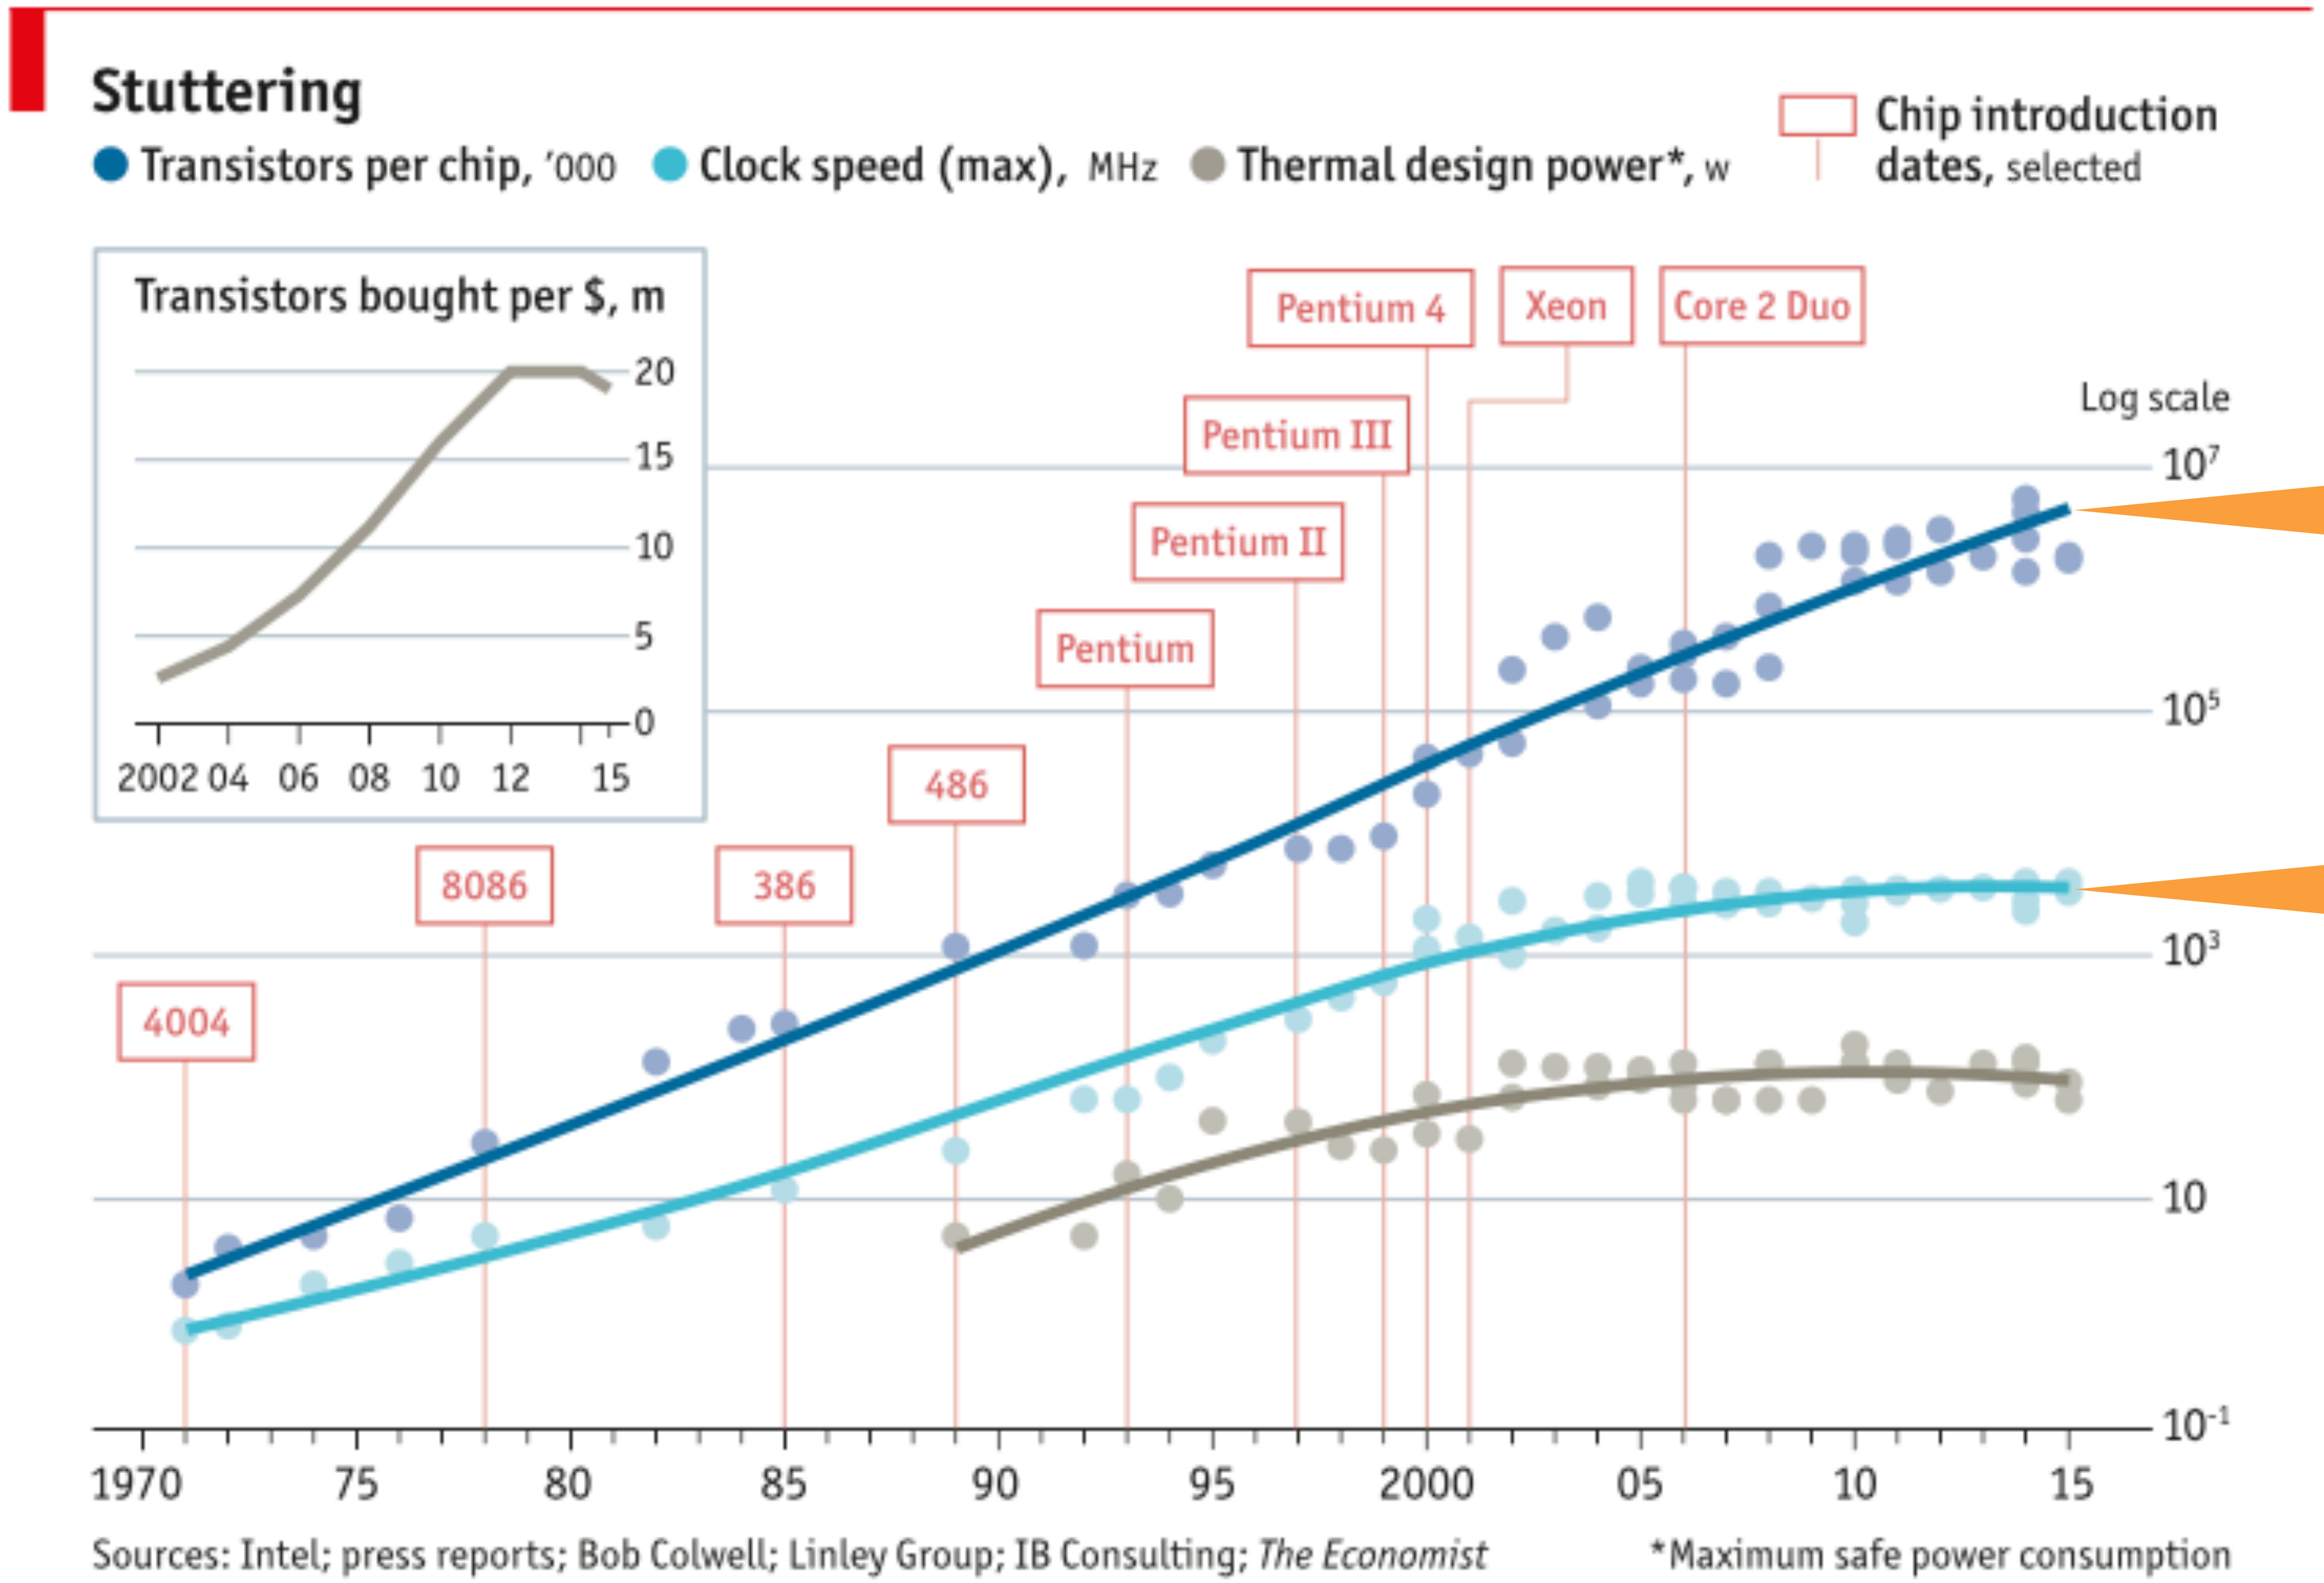

http://www.economist.com/technology-quarterly/2016-03-12/after-moores-law

Transistors per chip

Clock Speed [MHz]

64 cores
256 threads

GPUs

- Switch to multi-threading as big as Fortran to C++ transition

The Future: Machine Learning

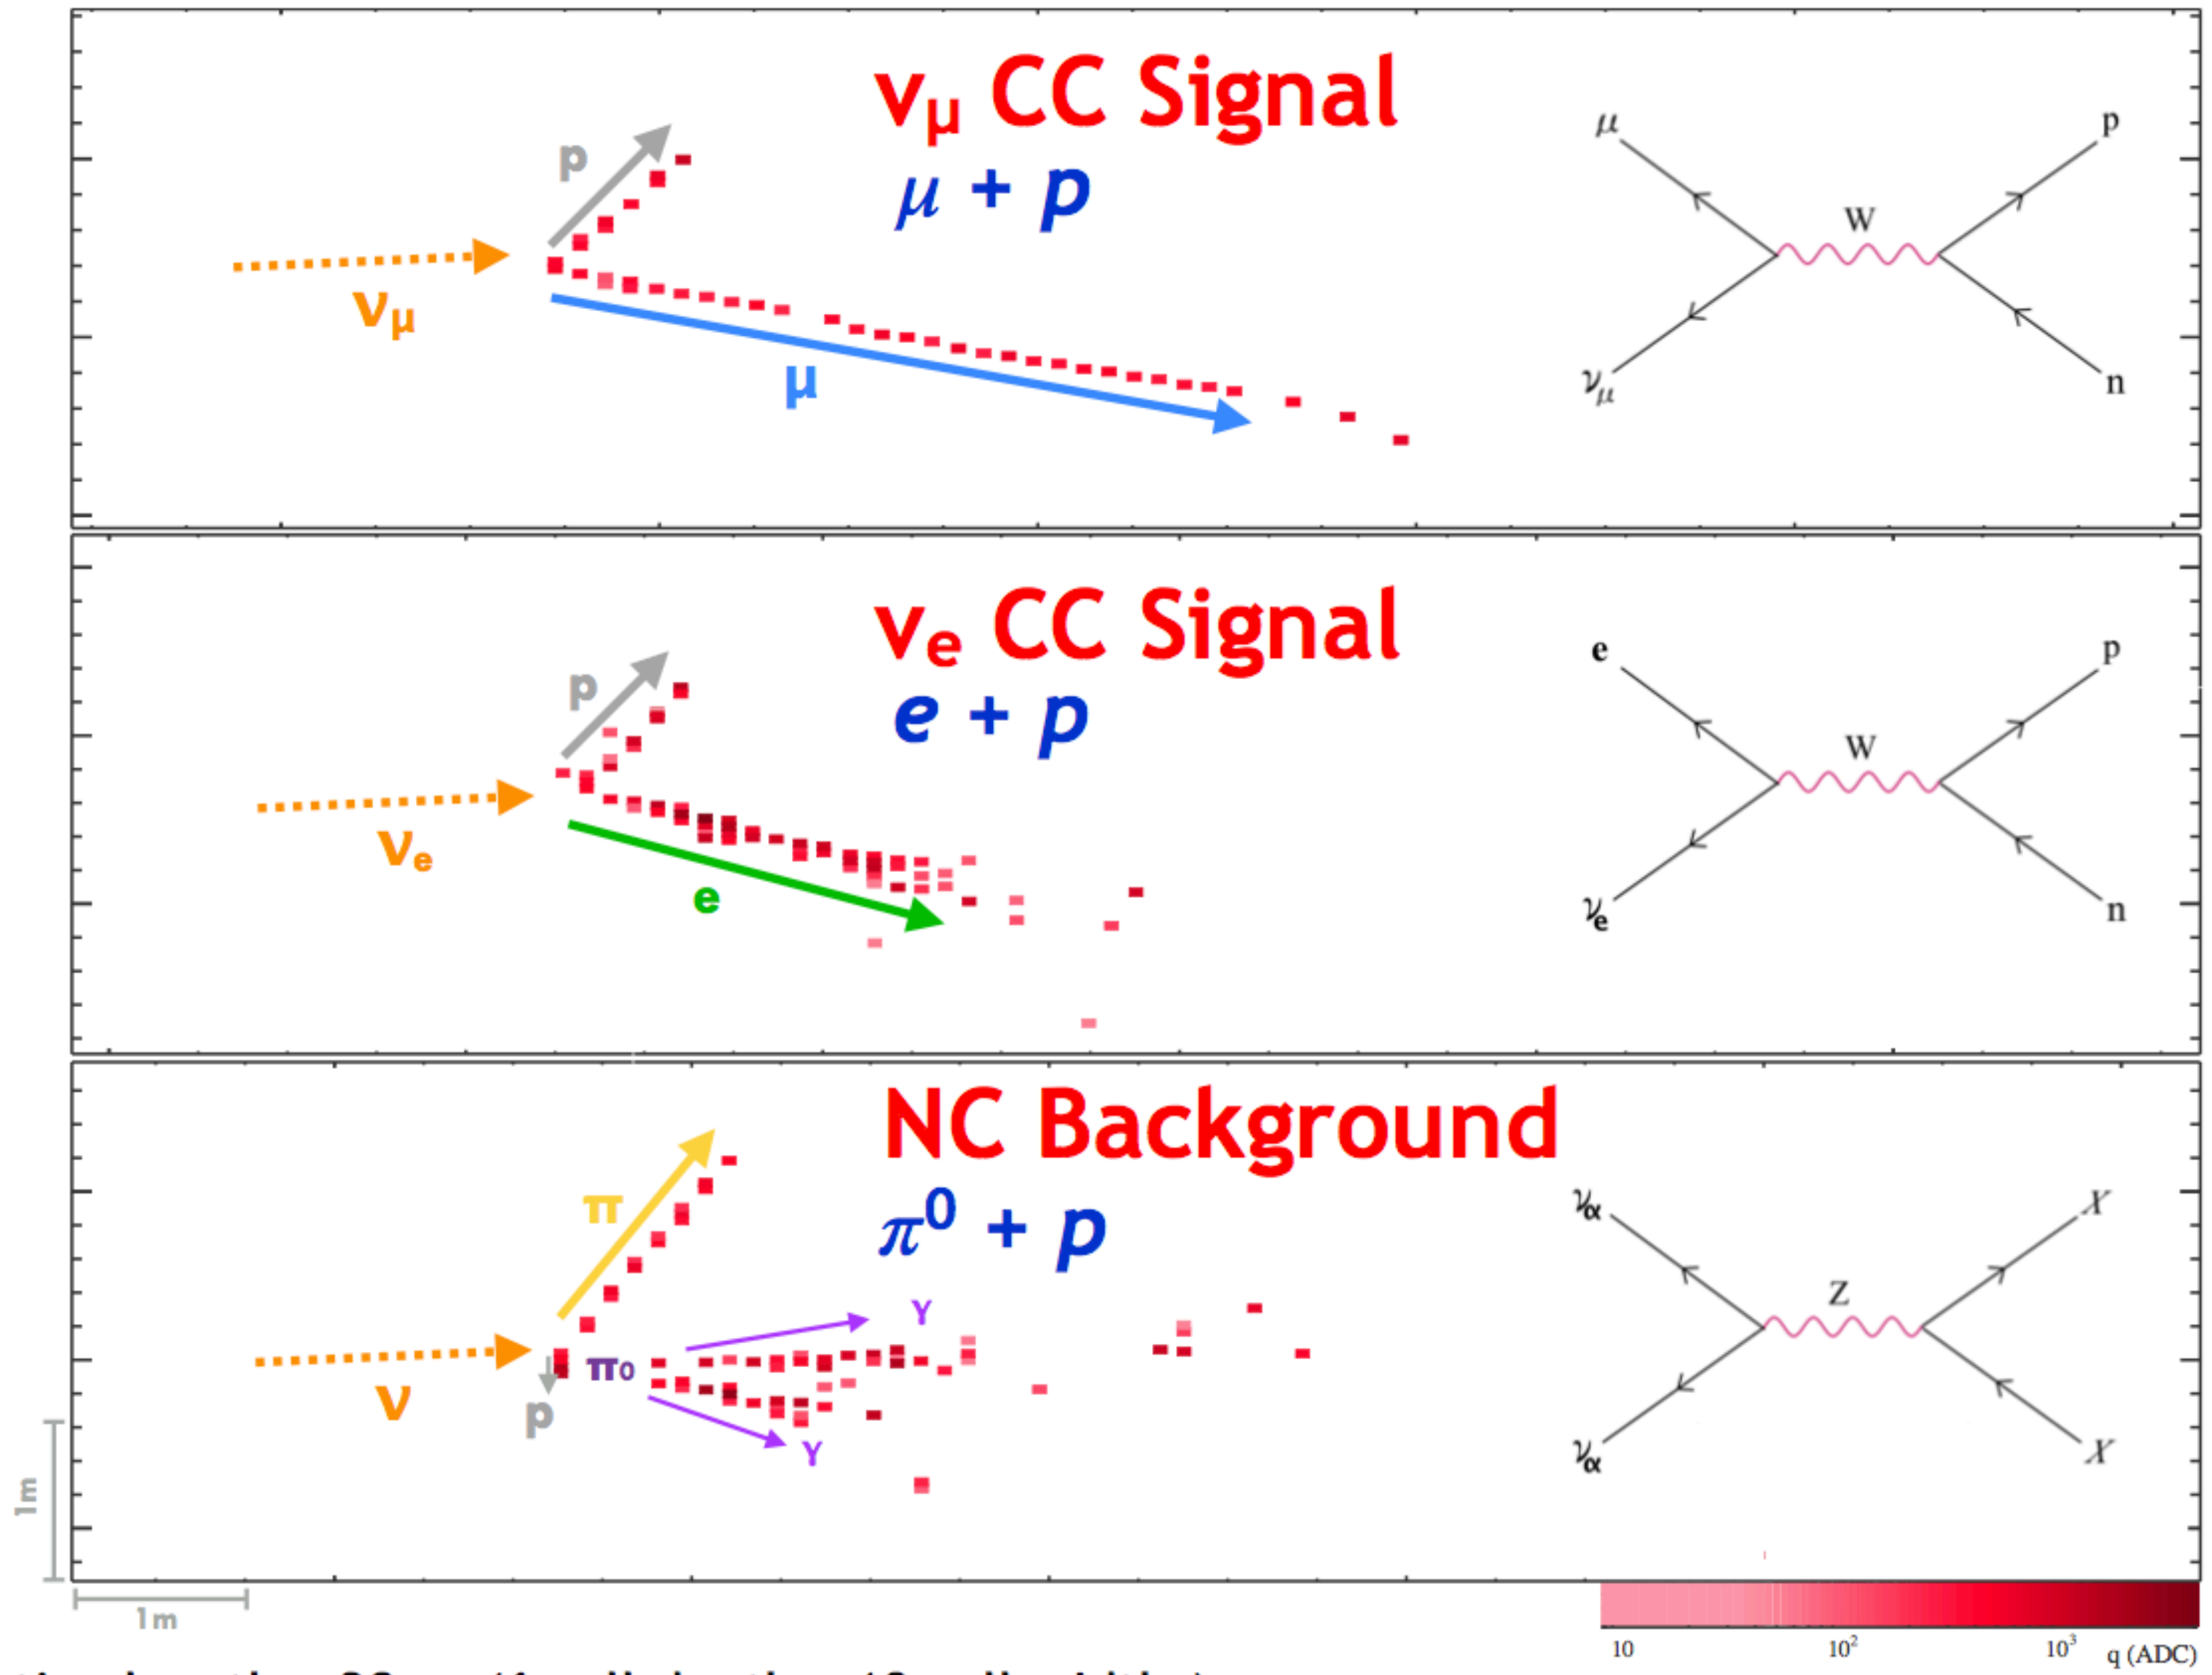

1 radiation length = 38cm (6 cell depths, 10 cell widths)

Example from NOvA (slides)

Presented at DS@HEP workshop at FNAL, May 2017

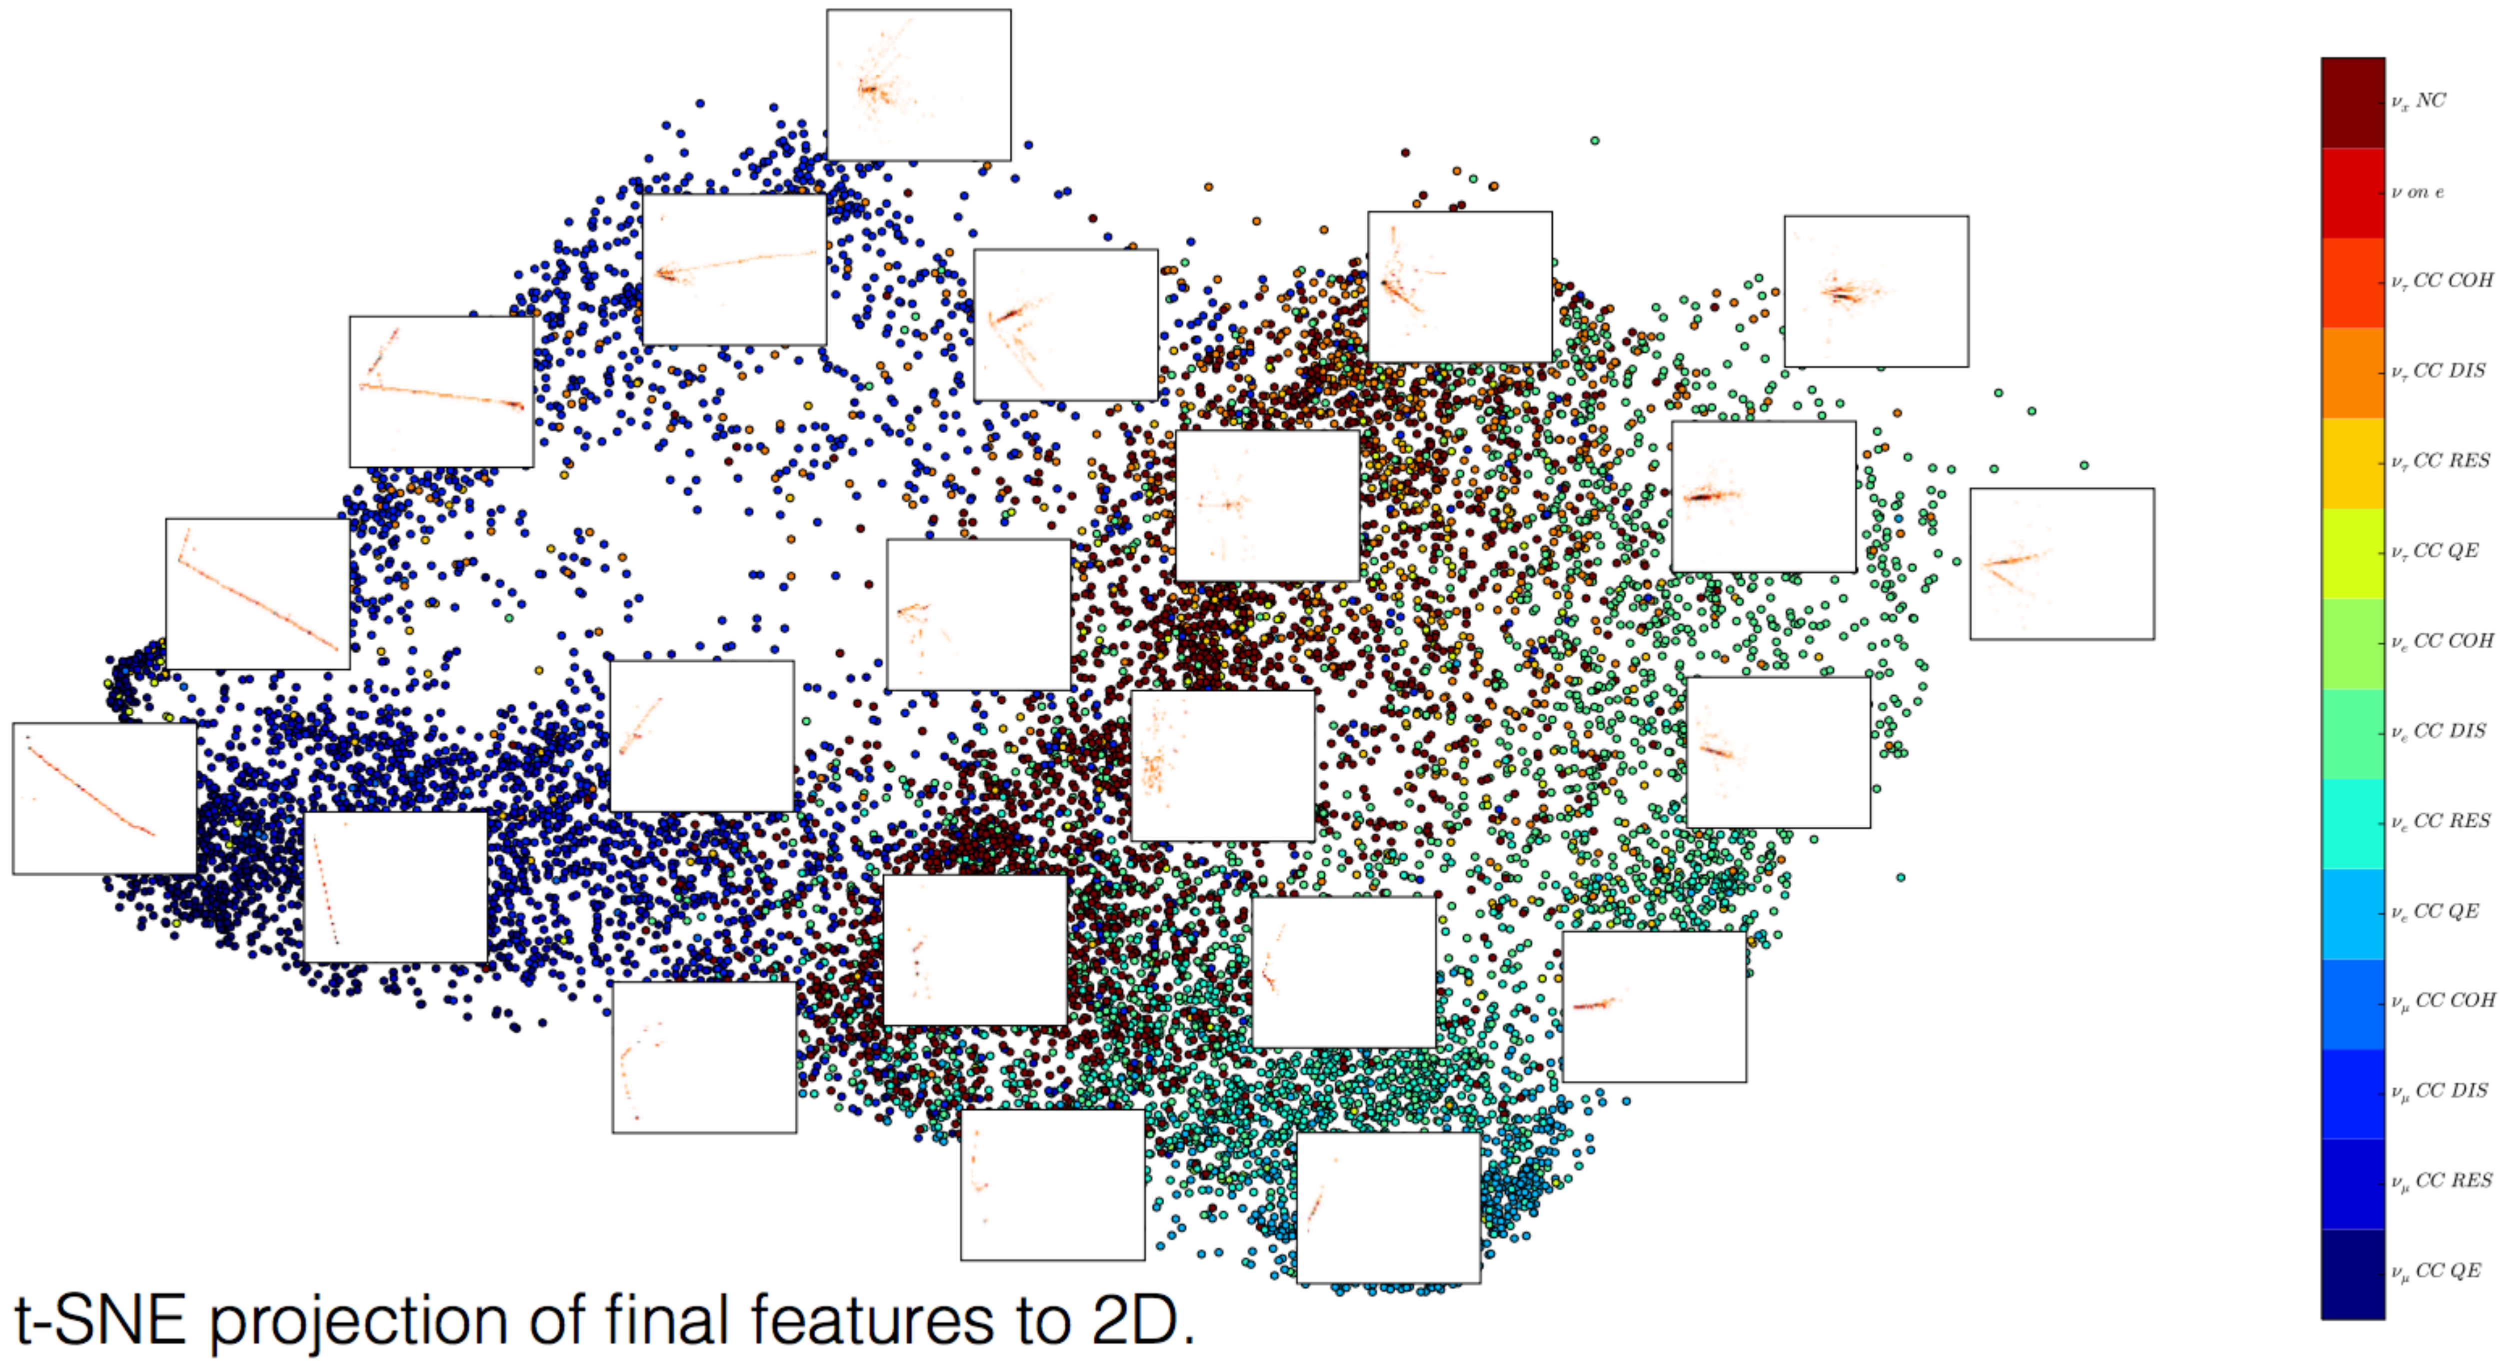

t-SNE projection of final features to 2D.
Truth labels, training sample subset.

- New methods and optimized hardware (GPUs, …) allow for deep learning networks

The Future: Data

Estimates of resource needs for HL-LHC

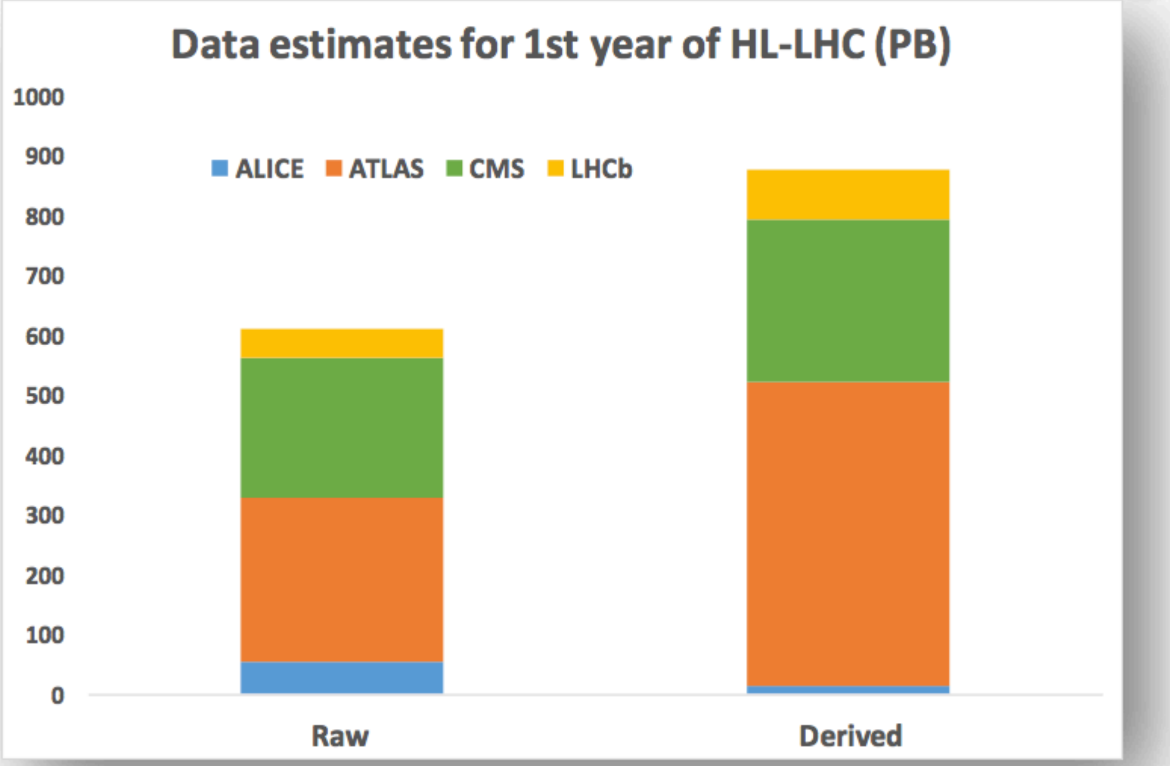

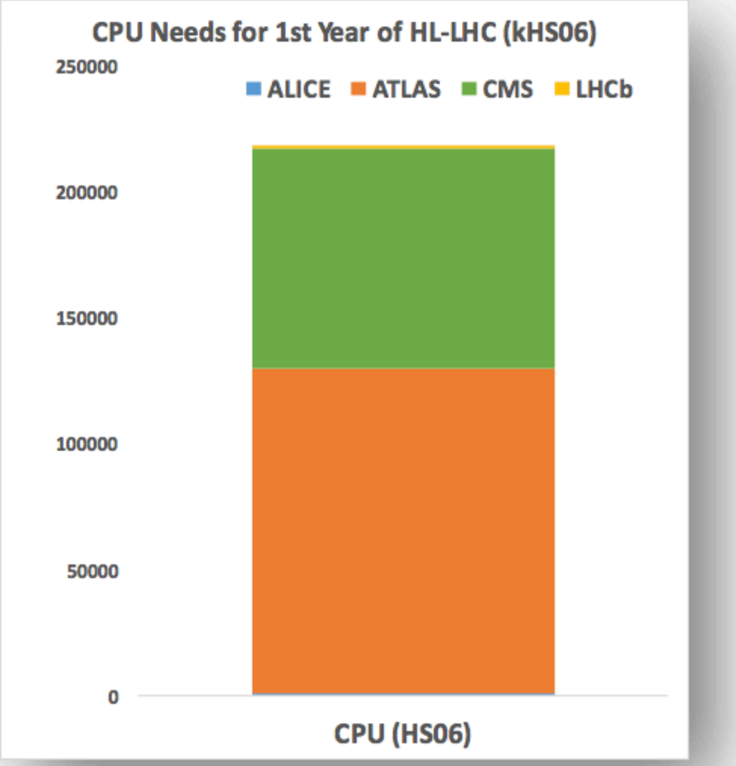

Data:
- Raw 2016: 50 PB → 2027: 600 PB
- Derived (1 copy): 2016: 80 PB → 2027: 900 PB

CPU:
- x60 from 2016

Technology at ~20%/year will bring x6-10 in 10-11 years

- Simple model based on today's computing models, but with expected HL-LHC operating parameters (pile-up, trigger rates, etc.)
- At least x10 above what is realistic to expect from technology with reasonably constant cost

CERN | WLCG 8 October 2016 Ian Bird 10

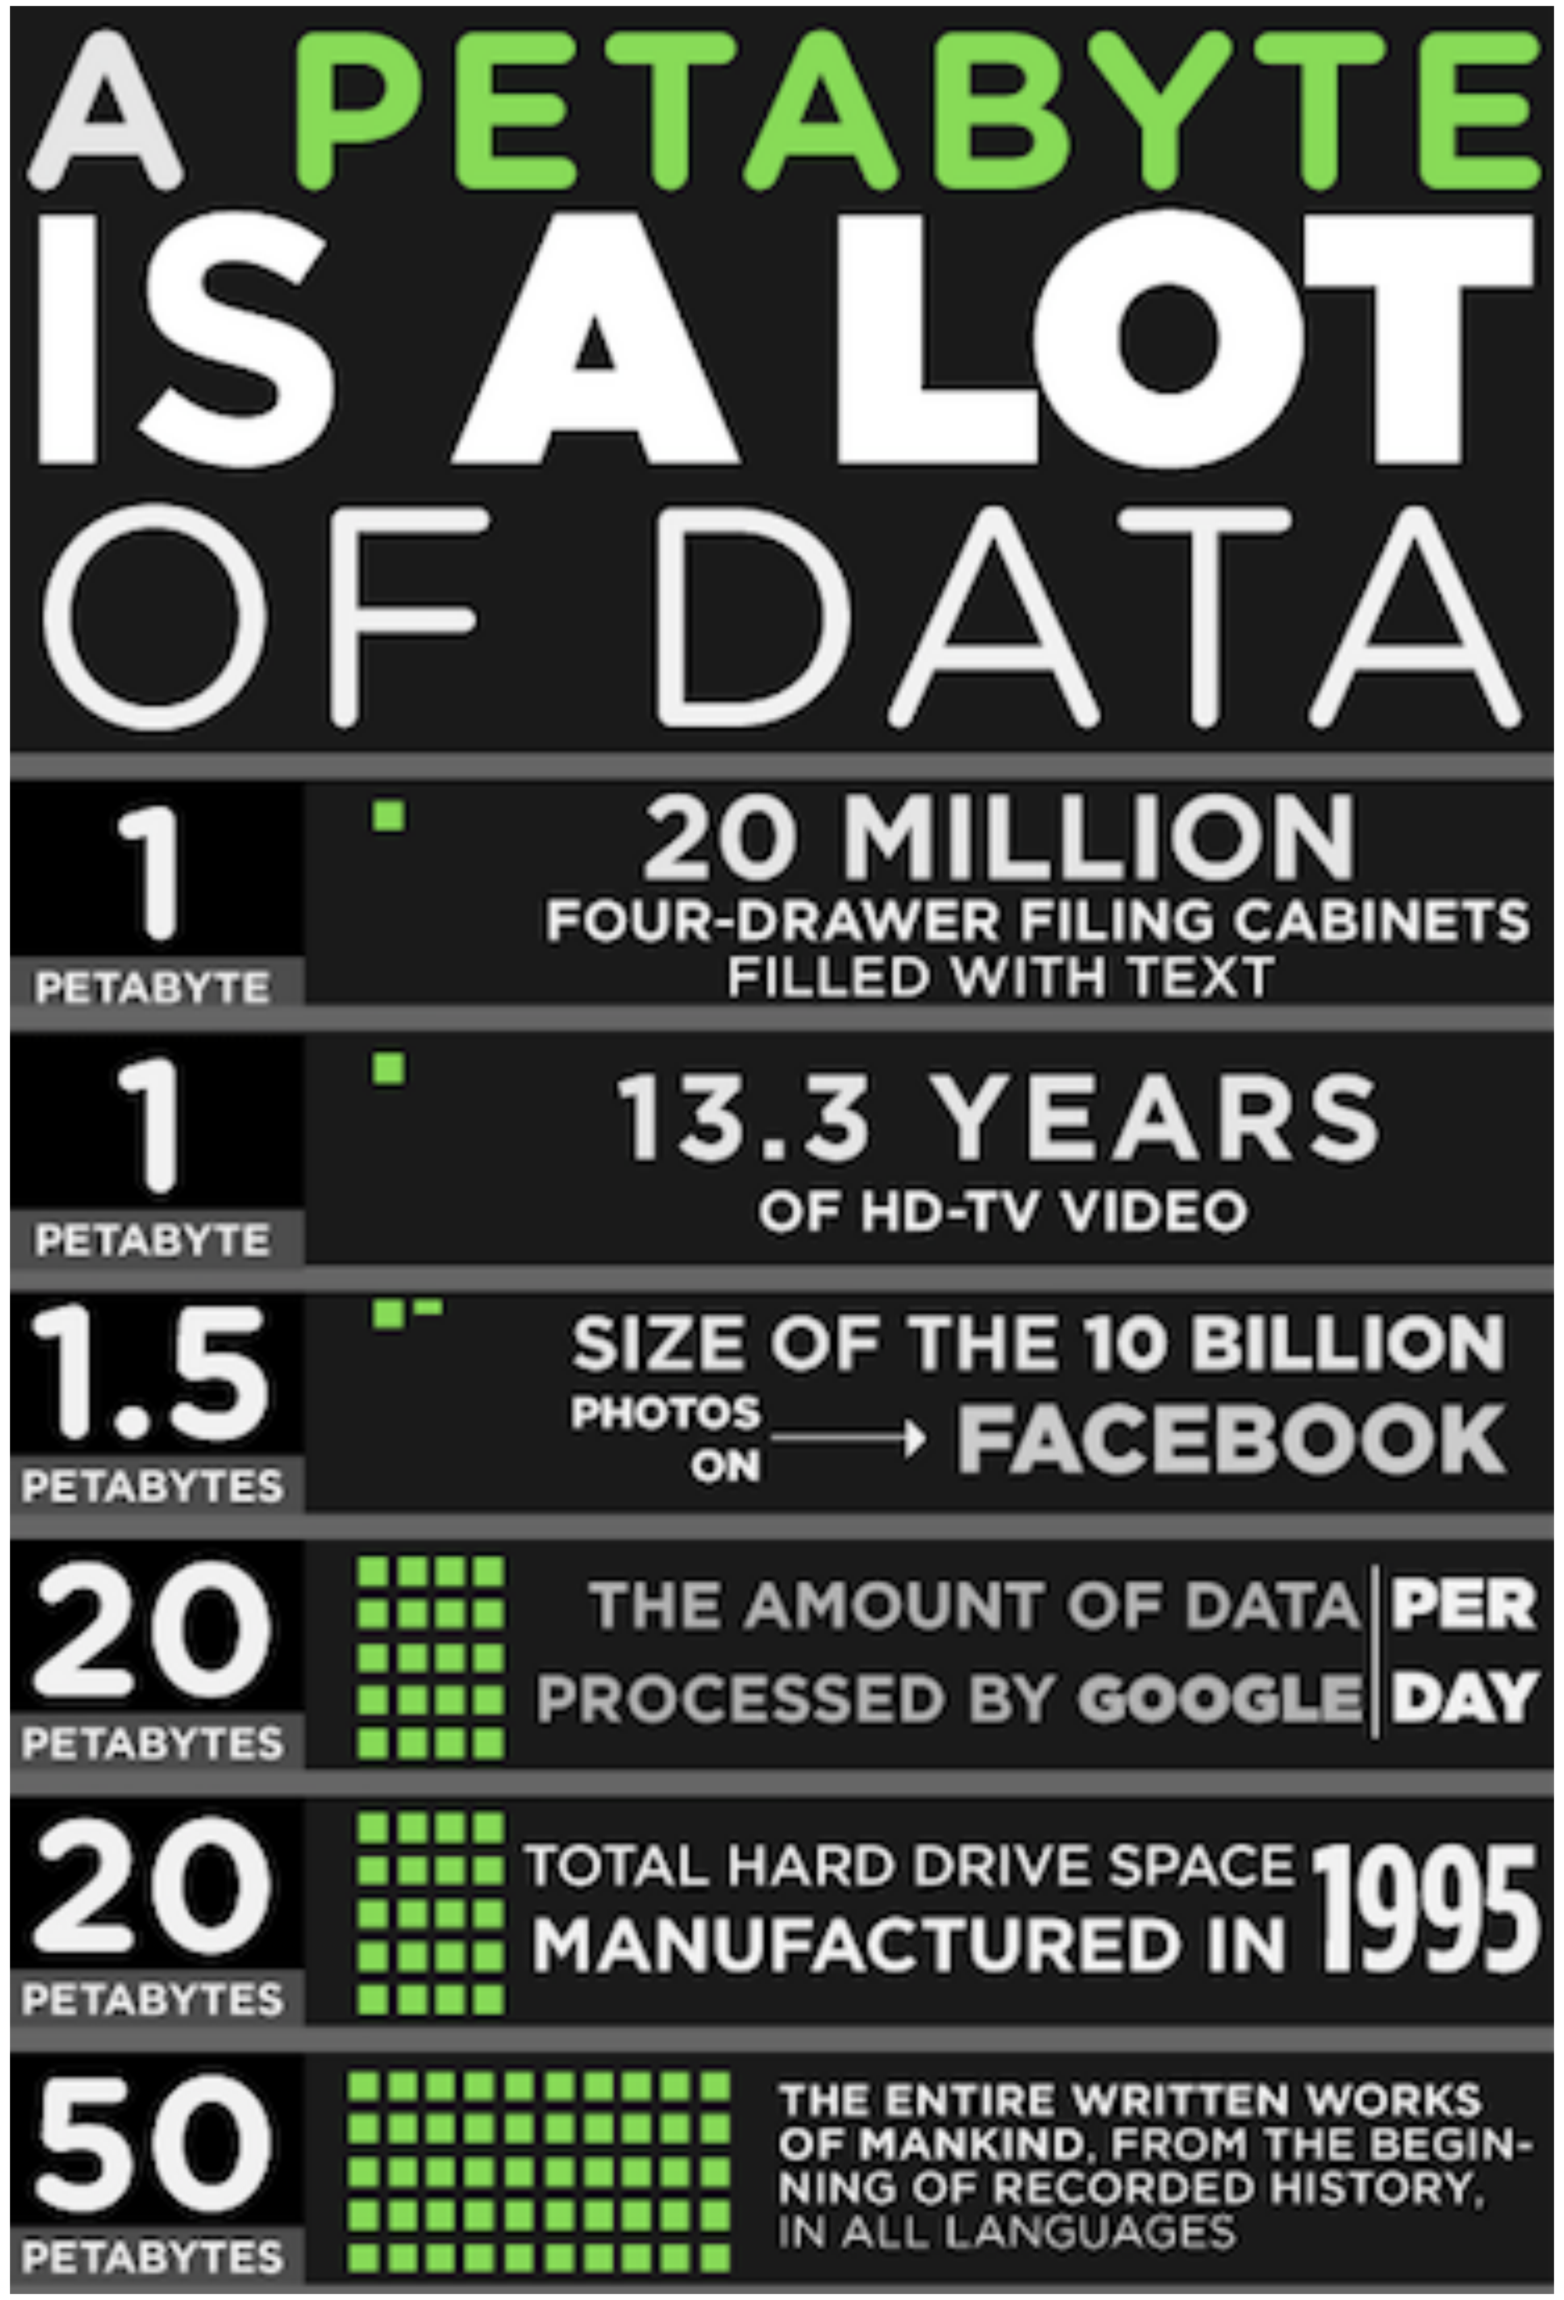

- Data volumes will increase to Exabyte scale after 2026 (HL-LHC and DUNE)

Fermilab

The Future: Processing Resources

Commercial clouds (Amazon Web Services, Google Cloud, …)

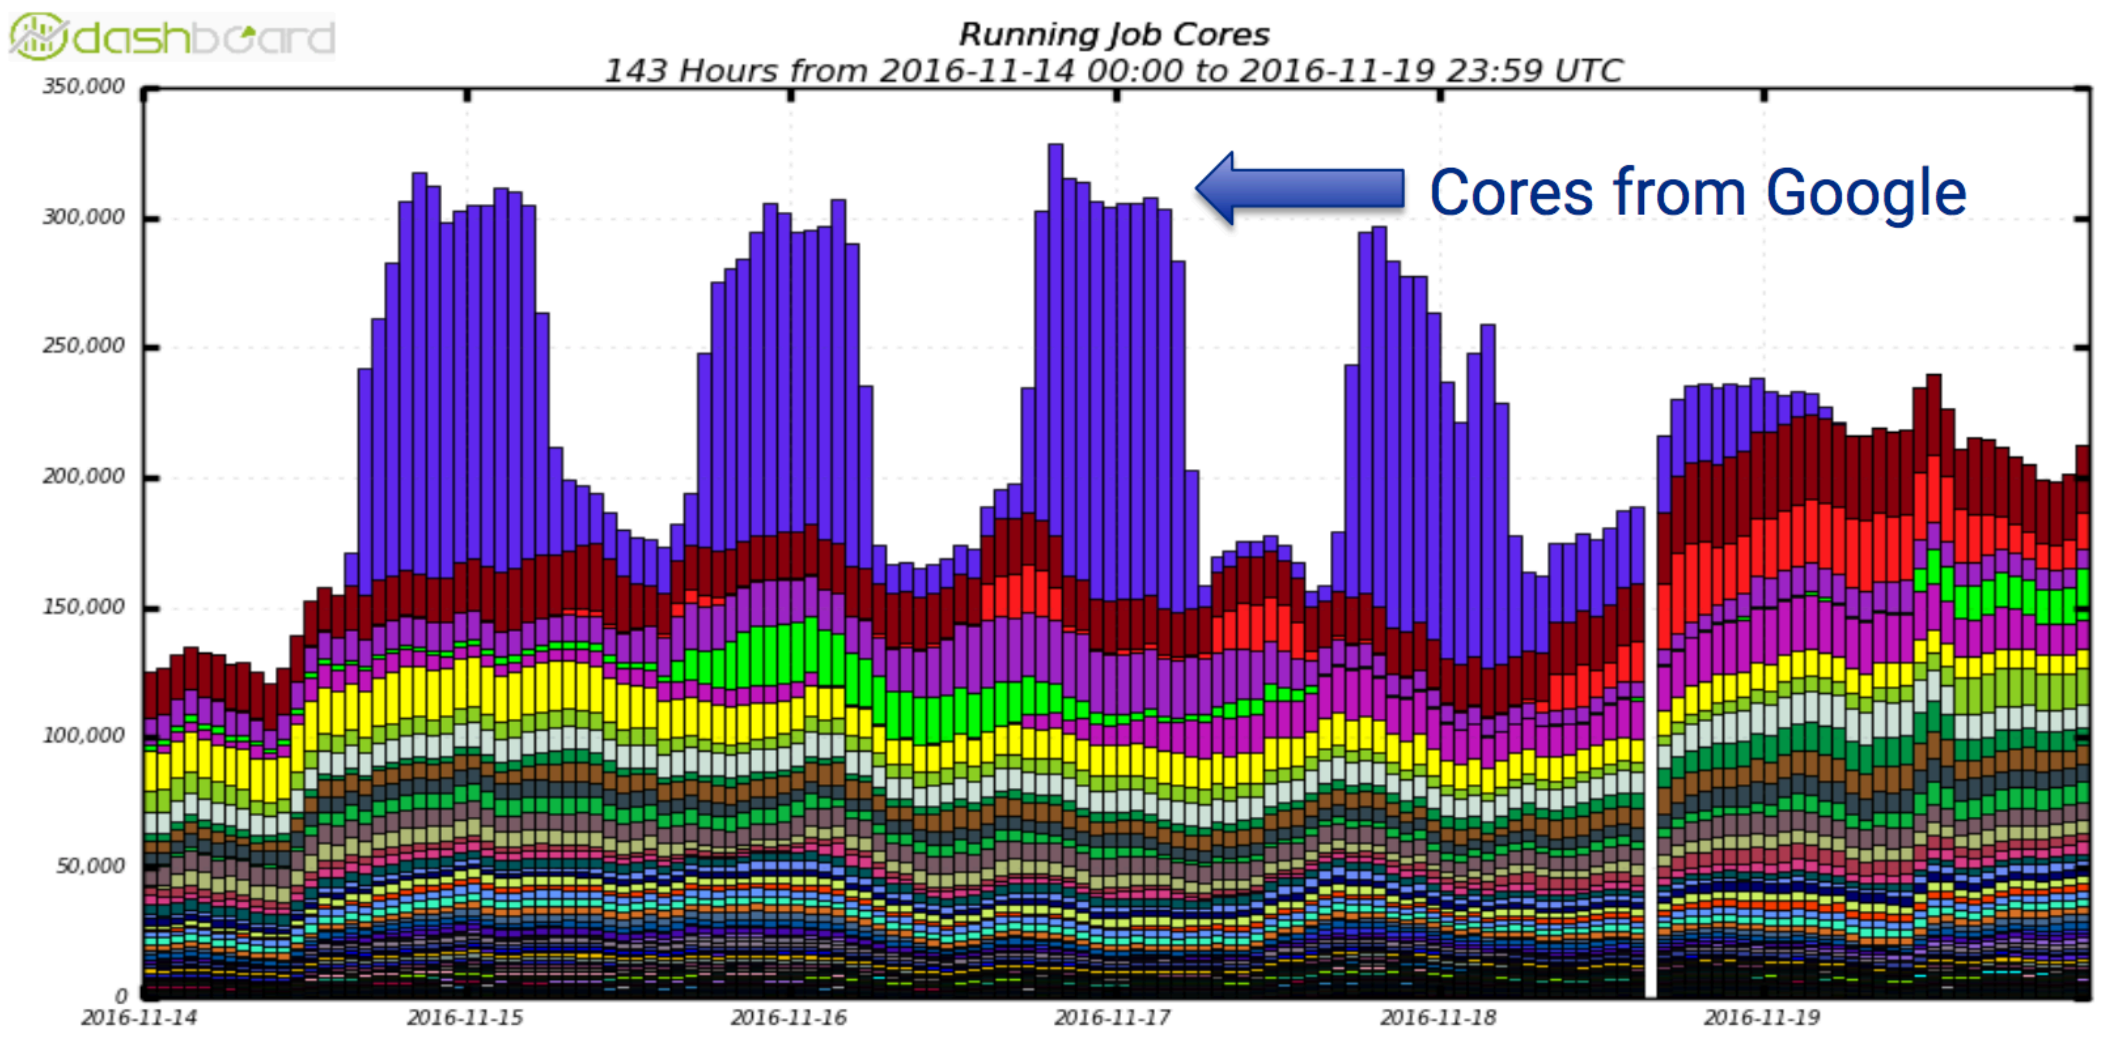

Exascale Supercomputers (HPC)

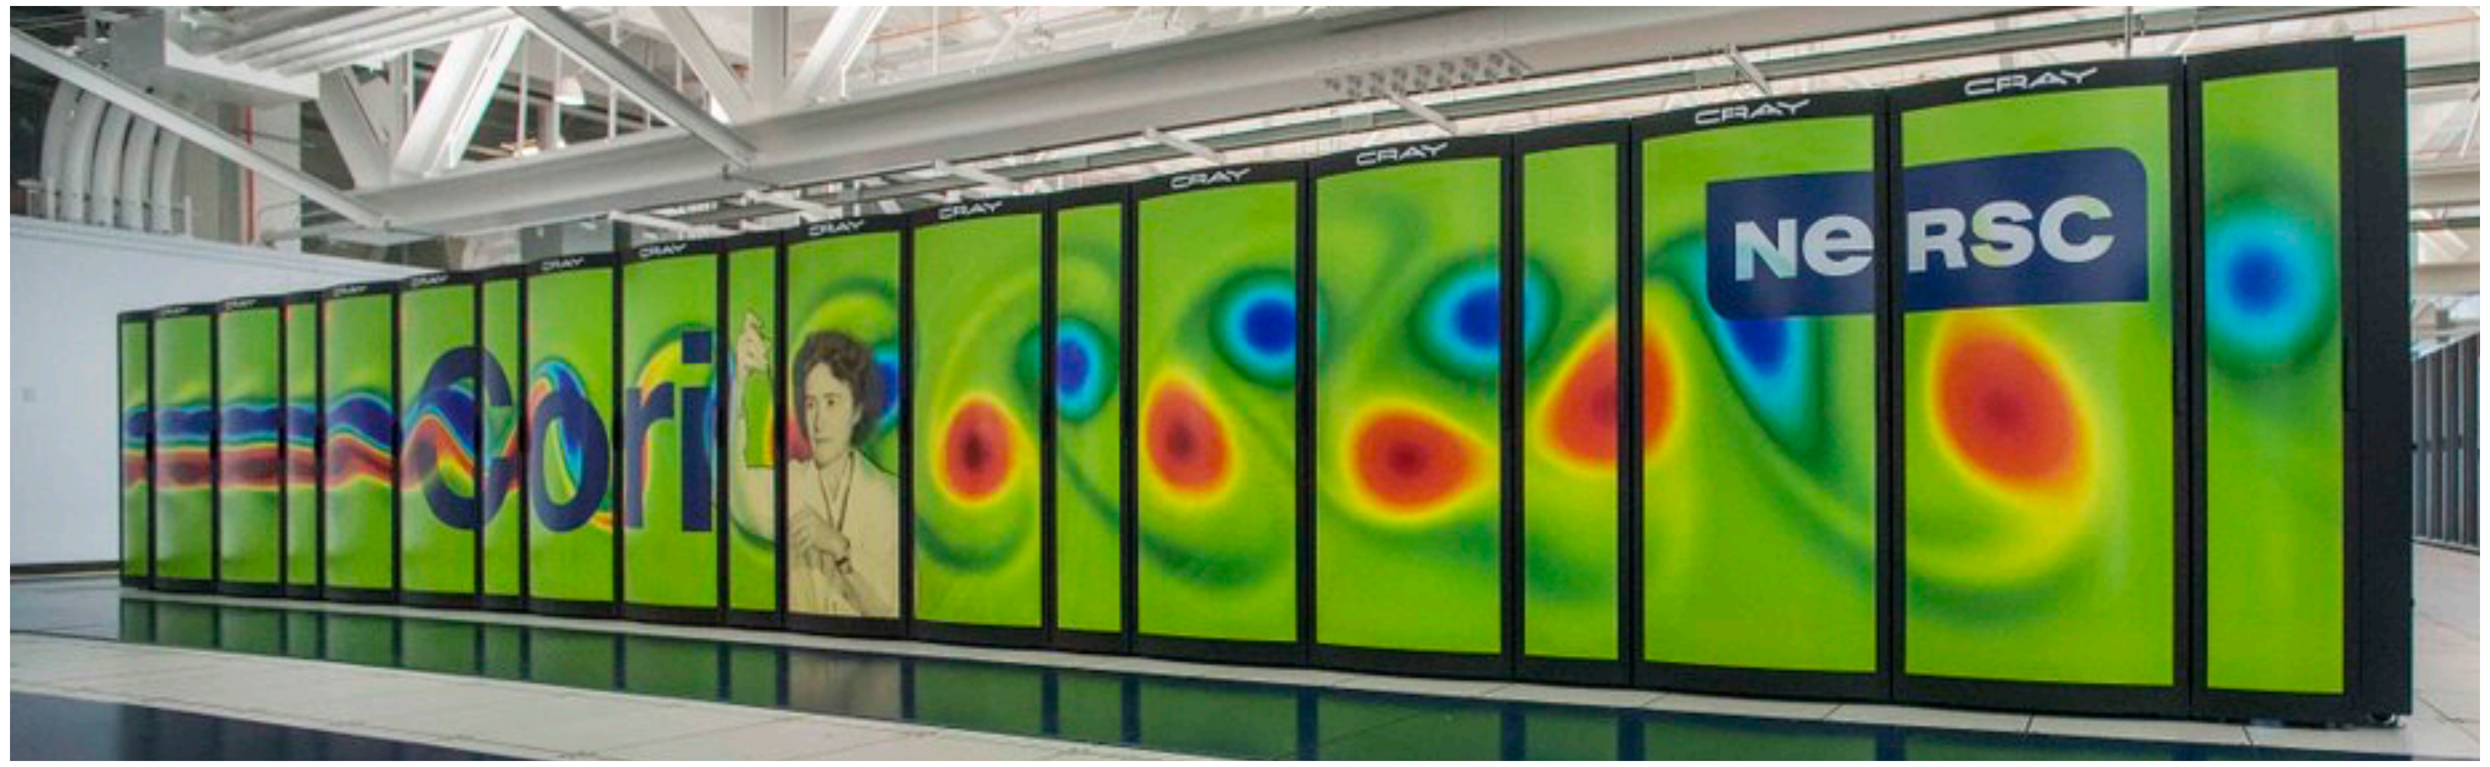

- The Grid will evolve into a more heterogeneous distributed system of Grid, commercial cloud and Exascale Supercomputers
- Community access to all forms of resources: HEPCloud

The Future: Quantum Computing

- Exploring overlaps between HEP and quantum computing
  - DOE roundtable Feb 2016
  - White House quantum computing meeting Oct 2016
  - DOE quantum testbeds workshop Feb 2017
  - Visits to Google and IBM
  - U. Chicago-Argonne-Fermilab quantum meeting May 2017

- Fermilab quantum initiatives
  - Quantum sensors
    - Adapting quantum devices for use as quantum sensors for particle physics experiments such as direct dark matter detection
  - Superconducting technologies
    - Some quantum computers use superconducting cavities similar to those we develop for accelerators…
  - HEP applications of near-term quantum computers
    - Google and IBM will make ~50 qubit quantum computers available for us to run on next year
  - Quantum networks
    - We have agreed to host a quantum network on site in collaboration with Caltech and AT&T

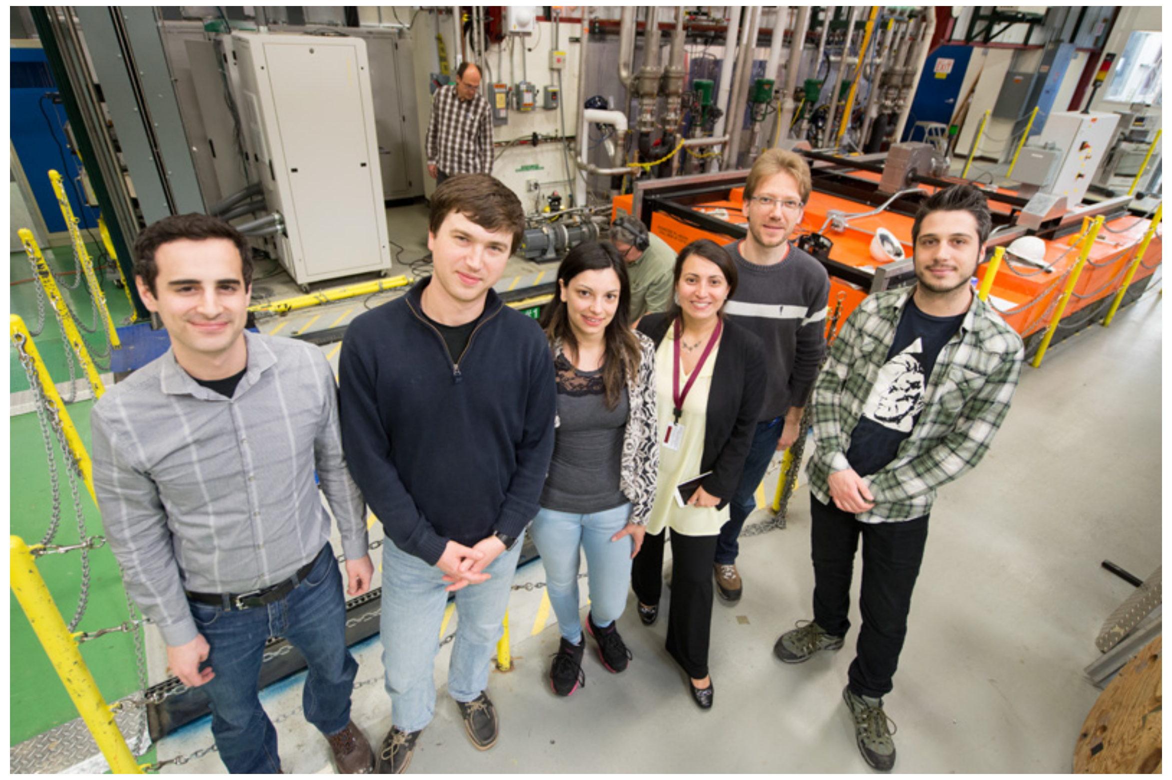

Conclusions

- Together we are stronger!
  - Large user community, large software developer base
  - Common solutions
  - Distributed but interconnected infrastructure, strong networks
- Agility is the key!
  - We started supporting many experiments in parallel
  - Followed by period where CDF and D0 dominated
  - Now going back to supporting many experiments
- For our science: we need maximum computing at minimum cost!
  - Use industry solutions where possible
  - But don't shy away from looking beyond what is available

- Fermilab was the nucleus for many innovations in the past!

- The future will be challenging but interesting and rewarding, and we will be right at the forefront of what is possible!

Thanks to the Scientific and Core Computing Divisions

Special Thanks to:

Jeff Appel, Tom Bozonelos, Liz Buckley-Geer, Joel Butler, Keith Coiley, Phil Demar, Dave Fagan, Irwin Gaines, Valerie Higgins, Mark Kaletka, Adam Lyon, Patty McBride, Hugh Montgomery, Gene Oleynik, Ruth Pordes, Liz Sexton-Kennedy, Marcia Teckenbrock, Rich Thompson, Marguerita Vittone, Margaret Votava, Vicky White, Steve Wolbers

BACKUP

Cyber retirement

1990

The emergence of the internet

1990

ROARK JOHNSON/BLACK STAR © 1990

Fermi Labs' Mark Kaletka and Philip Demar survey campus map.

Fermi Labs eases traffic crunch on campus DECnet

User installs separate LAN as 'firewall' to protect local network from wide-area routing congestion.

By Jim Brown
Senior Editor

BATAVIA, Ill. — Faced with runaway network growth, Fermi National Accelerator Laboratory has devised a local-area network routing backbone to better manage wide-area links and is laying plans to migrate to DECnet Phase V to solve addressing limitations in its current Phase IV net.

Fermi, a test center for physicists who conduct experiments using an on-site particle accelerator to collide subatomic particles, supports a DECnet with almost 20,000 devices. That network is surpassed in size only by Digital Equipment Corp.'s internal network.

Fermi's High Energy Physics network (HEPnet) enables researchers here and at various universities to access data and applications residing on the lab's DEC VAXes and Amdahl Corp. mainframes. HEPnet also lets researchers share findings with colleagues at other universities via electronic mail and file transfers.

Network growth is stretching the DECnet Phase IV installation to the maximum number of users that can be supported under the Phase IV addressing scheme, said

(continued on page 49)

Fermi Labs eases traffic crunch

continued from page 1

Philip Demar, Fermi's network analyst in charge of wide-area networking.

In fact, the lab has been rationing the few precious DECnet addresses it has left.

Consequently, Fermi expects to migrate to DECnet Phase V as

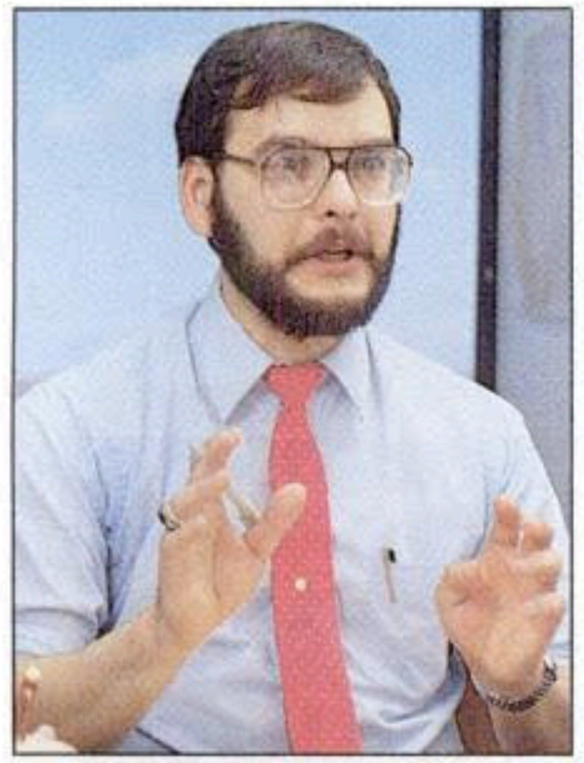

Mark Kaletka

soon as it is available and could begin beta-testing the software as early as next month. DECnet Phase V would solve Fermi's single greatest problem: accommodating growth on the network while resolving routing issues unaddressed in Phase IV.

"From our point of view, we desperately need Phase V of DECnet right now," Demar said. "We need it for the address space, and we also need it to control routing boundaries." DECnet Phase IV does not support boundaries, which enable users to make one physical network appear as several logical networks in order to avoid routing problems.

Sciences network (ESnet), Internet, the Magnetic Fusion Energy network and Bitnet, an electronic mail net. Additionally, the routing backbone links remote nodes to ESnet.

A cisco Systems router, which Fermi refers to as a firewall router, sits between the campuswide DECnet and the routing backbone, ensuring that packets routed between remote sites do not congest the campus DECnet. The firewall router routes packets from the campus DECnet to the router that provides access to the appropriate WAN.

Initially, HEPnet was quite simple and did not generate enough traffic to congest the campuswide DECnet. Five years ago, HEPnet consisted of five star-configured WANs linked together by DECnet routers and 9.6K bit/sec circuits. The hubs of all star networks were located at

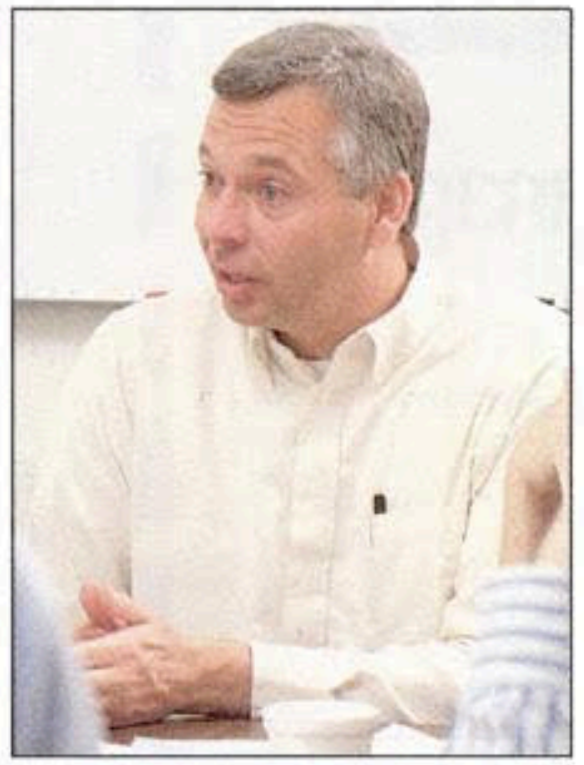

Philip Demar

labs that supported point-to-

HEPnet (ESnet-DECnet) Topology

(as of May, 1990)

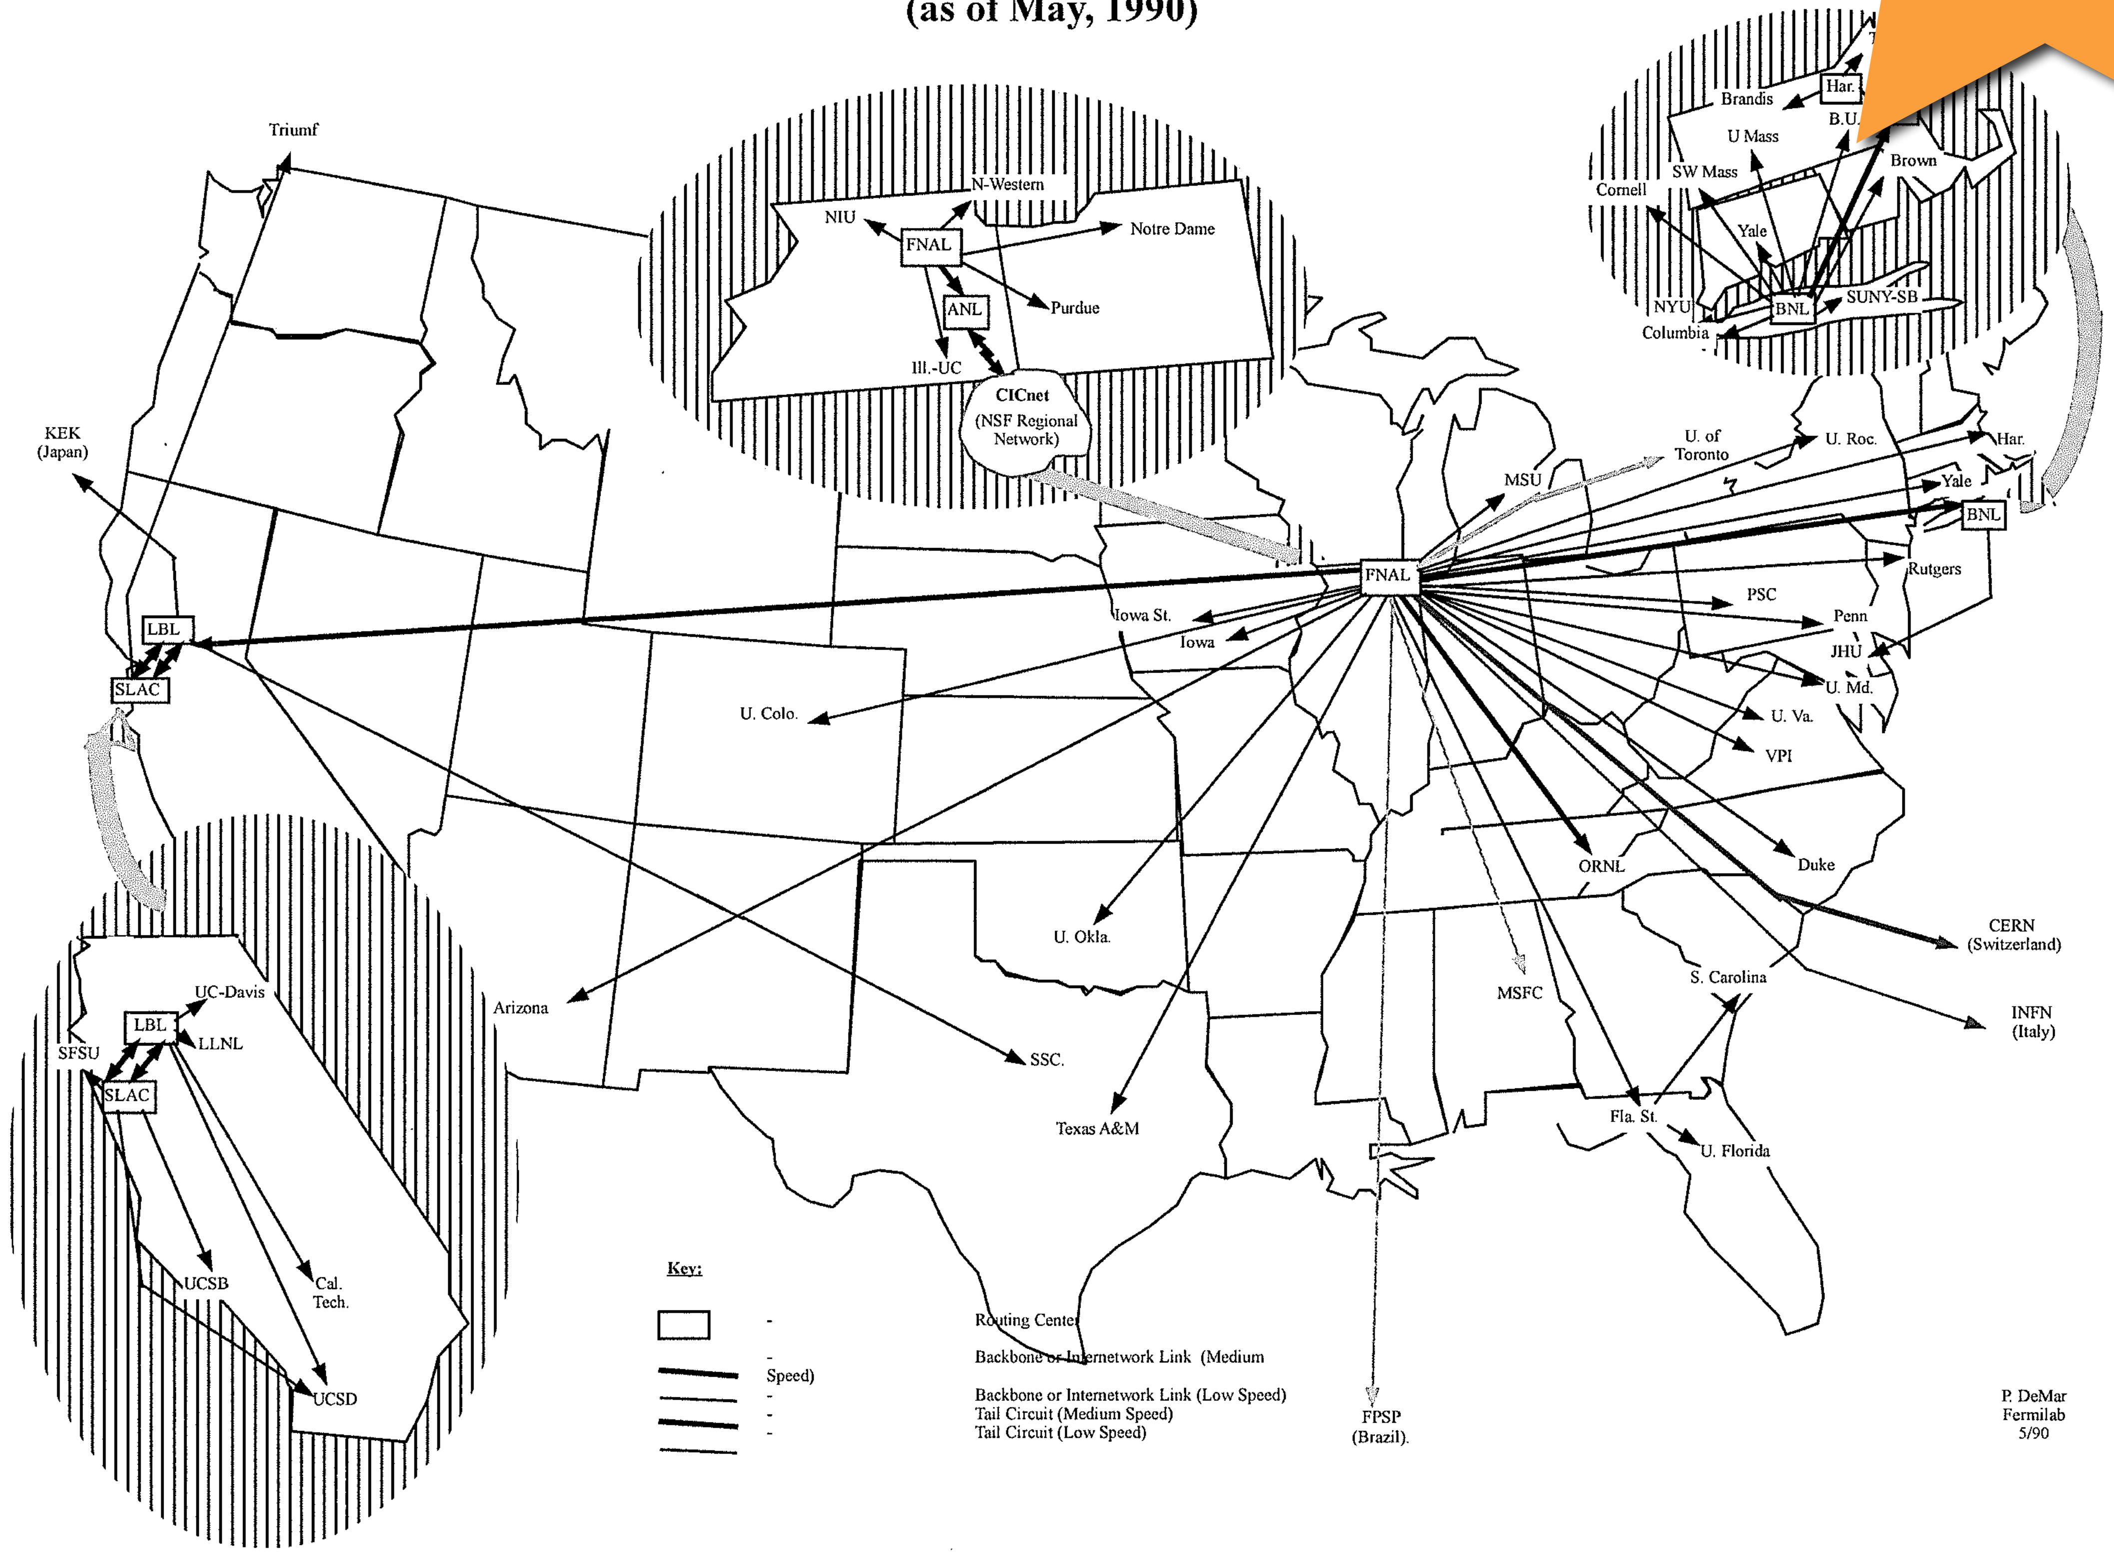

- ESNet, supporting both DECNet and the internet (TCP/IP)
    - The internet won!

Fermilab

Tim Berners-Lee at FNAL in the early days of the web

1991

FermiNews
Fermi National Accelerator Laboratory
Friday, August 16, 1996 Number 16

INSIDE

Detour Ahead

East Entrance to Close Temporarily

Fermilab will close the east entrance to the Laboratory from 6 a.m. Monday, August 19 through 8 p.m. Friday, August 23, to permit maintenance on the railroad grade crossing by the Elgin, Joliet and Eastern Railroad. During the temporary closing, employees and others who normally enter via Batavia Road on the east side of the Laboratory must enter instead via Pine Street on the west. The entrance will also be closed to bicycles and pedestrians.

Detour signs will reroute traffic via Route 59 south, Butterfield Road west, and Kirk Road north to Pine Street. Questions? Call the Office of Public Affairs at 630-840-3351.

Photo by Jenny Mullins

Glenn Blanford, a graduate student from the University of California-Irvine, at work on the new Fermilab home page that will debut at the end of August.

High-Energy Physics

Birthplace of the Web

by Eric Berger, Office of Public Affairs

In mid-1991, when it spun its first thread from CERN to Fermilab, hardly anyone had ever heard of the World Wide Web. Now, only four years later, the Web is woven through the fabric of our culture, connecting people all over the planet in a new medium of communication. But of the millions who type "http:/ / www...." each day, how many realize that—like the universe itself—it all began with high-energy physics?

What started out as tool for far-flung scientific colleagues to share each others' data may ultimately count among high-energy physics' most significant contributions to modern technology.

The Web was born in 1990, when Tim Berners-Lee, a computer scientist at CERN, the European Laboratory for Particle Physics, programmed the first types of computer codes, called protocols, that allowed a computer any-

continued on page 8

The Web at Fermilab

continued from page 1

Tim Berners-Lee, who wrote the software protocols to set up the World Wide Web in 1990 and 1991. He now directs the W3 Consortium, an open forum of companies and organizations with the mission to realize the full potential of the Web.

where to contact any other computer and communicate unfettered by log-in accounts and database incompatibilities.

"He's someone who thinks very clearly about the fundamentals, but besides this, he happened to come upon his idea at the right time," said Ruth Pordes, Computing Division Online Systems Department head. "His idea was to make distributed information over a network easily accessible to the whole world, and to work out the protocol of how to do that."

In its infancy, an internal laboratory version of the Web allowed high-energy physicists at CERN to share data and papers within their own network, but it remained difficult to connect to outside servers—from CERN to Fermilab for example—to share documents with other scientists.

Pordes, who had met Berners-Lee at international meetings, invited him to visit Fermilab for a few weeks in January, 1991. The visit coincided with a conference on hypertext, familiar to most Internet users as the "link" function of Web browsers. The conference gave Berners-Lee a stimulus for developing his ideas.

Berners-Lee returned to Fermilab in 1992. During his visit he made the first one-click link between the CERN server and Fermilab's central computers.

"We installed it (the CERN Web server) on FNALV and we linked it up to the documents used for data-taking and later analyses," Pordes said.

Fermilab's Jonathan Streets, leader of the Experiments Online Support Group, helped Berners-Lee set up the system.

"He and I wrote the server on FNALV that served the documents," Streets said. "Now anybody could come in and get them. That was the first time anybody could use the same interface to read documents pertaining to both data-taking and analyses—because we had very different databases."

Before Berners-Lee wrote the hypertext transfer protocols each computer server generally required a separate log-in account and had a different documentation system, making smooth communication between systems difficult. Berners-Lee's software (the Web), broke down the barriers dividing the hardware (the Internet).

The following summer, Berners-Lee gave the first United States presentation on the World Wide Web's merits.

"Ruth basically has to take credit for spotting the Web," said Vicky White, Computing Division deputy head. "We both knew Tim Berners-Lee's work, but she showed the first interest in it."

That Was Then...

Berners-Lee used these two transparencies in a 1992 Fermilab talk highlighting the World Wide Web.

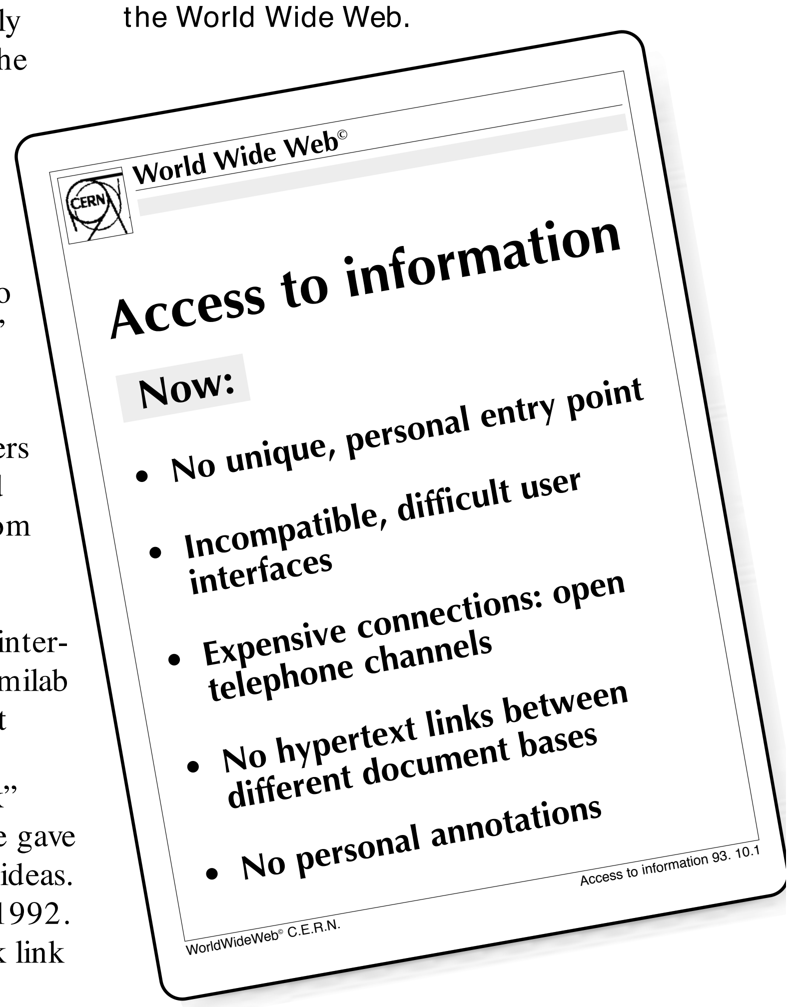

Fermilab joined the World Wide Web early. The Sloan Digital Sky Survey, an ambitious sky-mapping project in which Fermilab collaborates, became one of the first projects to adopt the Web in 1992.

"Sloan is a project that has a far-flung collaboration, and we felt having information that could be easily retrieved would be beneficial," Pordes said.

Beyond High-Energy Physics

Berners-Lee and others soon discovered the Web applied to areas well beyond high-energy physics.

"After that, he started working on protocols for the World Wide Web in order to have communication not only among physicists, but also to be able to find information from all the various other places available," Pordes said.

The next step required making the technology more accessible. Current Internet users surf the web using browsers such as Netscape or Web Explorer that offer windows into the Web. But in 1992 none of these mouse-friendly programs existed. A program called Mosaic paved the way for later browsers.

"Around 1992 NCSA (National Center for Supercomputing Applications at the University of Illinois at Urbana/ Champaign) started work on Mosaic," Pordes said. "In 1993 I invited Tim to Fermilab. I contacted NCSA and sug-

- "That was the first time anybody could use the same inter- face to read documents pertaining to both data-taking and analyses—because we had very different databases."

Fermilab

Mainframes vs. Farms

- First Unix Farms (IBM & SGI) vs. SGI mainframe

Disk Caching

- dCache reaches 1.7 PB

Tape capacity increases

- Old robots get dismantled, new robots get installed
- Total capacity now close to **1 Exabyte**