

Computing needs for running experiments

- Look at 2 experiments:

	BaBar	CDF
Start of data taking	1999	2002
Run until	2010	2010
Italian physicists	~80	~60
- CDF is an evolution of previous experiment (1985)
- BaBar is brand new
- How much data/CPU ?
- What is done in US, what in Italy ?
- What are the needs in the INFN sites ?
 - Skip national farms (Caspur, CNAF)
- **Their data, my opinions**

Centralita' dei servizi di calcolo: oltre il "computing"

- 90% dell'attivita' di un ricercatore e' davanti ad un computer
 - Posta elettronica
 - Lettura documenti/articoli (biblioteca → web)
 - Sviluppo software
 - Analisi
 - Scrittura documenti, presentazioni, articoli
 - CAD (elettronico, meccanico)
 - Test di elettronica/rivelatori
 - Controllo esperimento (anche remoto)
- Tutto questo dipende dai servizi di calcolo della sezione
 - GRAZIE

29/6/99 Commissione Calcolo a Napoli: Bisogni di CDF nel Run II (2000 - 2003 - ...)

- Cominciamo dalla fine, i bisogni:

- RETE! RETE!! RETE !!!

- intercontinentale
- italiana
- locale

OK! OK!! OK!!!

Ma: niente relax

- system management: linux ~ci siamo, ancora ~nascosto
- licenze: compilatore C++: KAI Niente licenza, gcc in arrivo
- supporto software: Root, C++, Java, SMQL Ancora li' + Python
- collaborazione a distanza: videoconf + application sharing
- ancora rete: Qualita' di Servizio

Ancora un bisogno, minori
speranze: bastano reti
sovradimensionate !?

CDF vs BaBar (in numbers)

- Data (Root I/O)
 - 10^9 events/year
 - 250KB/event
 - 500TB/year data (not yet)
- HW at FNAL
 - 200 duals now (1K coming)
 - 128-CPU SMP leaving
 - 30TB disk (200 coming)
- User's analysis
 - ~100 data sets ~1TB each
 - TOO MUCH TO COPY
 - ☞ CENTRAL MODEL
 - ☞ All at FNAL
- Data (Objectivity + Root I/O)
 - $5 \cdot 10^9$ events/year
 - 30KB/event
 - 550TB total data now
- HW at SLAC
 - $O(1500)$ machines now
 - 40TB disk
- User's analysis
 - "micro" now ~10TB total
 - CAN COPY
 - ☞ HIERARCHICAL
 - ☞ Tier A-B-C

CDF vs BaBar (in words)

- Similar disk/CPU requirements
 - Central analysis farm: $O(1000)$ machines
 - Central data repository: $O(100)$ TB on disk, $O(1$ PB) on tape
 - ☞ CDF has more (more complex) data
 - ☞ BaBar has Objectivity
- CDF started with one big + many very small computing centers
 - can't move data
 - political/practical reasons (manpower e.g.)
- BaBar started with one big + few medium + many small
 - can move data
 - Objectivity encourage this (but makes it very hard)
 - political/practical reasons (cost share)

Computing Model Dictionary

- Try to use common definitions (from BaBar)
 - **Tier A** ~1000 CPU
 - ☞ stores raw data (the PetaBytes)
 - ☞ run reconstruction
 - ☞ allow analysis of "ALL" reconstructed data
 - ☞ produces MC
 - ☞ accessed from everybody, anytime (7x24)
 - ☞ provides several GRID-like services
 - **Tier B** ~ 100 CPU
 - ☞ allow analysis of "SOME" reconstructed data
 - ☞ produces MC
 - **Tier C** ~ 10 CPU
 - ☞ allow analysis of "LITTLE" reconstructed data

CDF/BaBar Italy Computing Models

- **Past** (plan)

- **Tier A/B** at FNAL
- **Tier C** in Italy (3~4)

- **Present**

- Building our **Tier B** at FNAL (easy, just buy)
- No **Tier C**

- **Next**

- **Tier B+** at CNAF
- GRID (EDG, DataTAG)

- **Future**

- Interactive at CNAF ?

- **Past**

- **Tier A** at SLAC
- national INFN **Tier B** at Caspur
☞ much, much work

- **Present**

- **Tier A** at Padova-SLAC-IN2P3-PPARC
- **MC** at Roma1

- **Next**

- GRID (EDG)

- **Future**

- **Tier A, MC**: Pd,Rm → CNAF ?

Computing Model evolution

- In CDF some "small" (Italy, UK...) growing to medium+
 - ☞ $O(100)$ nodes, access to all collab (no Tier spec)
 - ☞ now can move lots of data
 - ☞ want share hw across countries (too many data)
- In BaBar top big one is splitting, therefore Italian medium one is growing (30→110 nodes)
 - BaBar keeps very distributed MC production
 - ☞ little impact on Italy
- CDF+BaBar converging toward one large computing facility (at CNAF) embedded in a world-wide GRID structure
- Both will do without local mini-farms
 - Taking data and analysing them is a lot of work
 - A large working farm is a lot of work
 - ☞ BaBar Padova: 12 people to build reprocessing farm

The BaBar reprocessing farm:

largest INFN computing facility

largest italian IBM Linux cluster

BaBar Reprocessing farm fully cabled

CDF needs

- Access to Fermilab
 - Kerberos. Soon add GRID (non-EDG)
- Access to Regional Center(s) (also outside Italy)
 - GRID
- Interactive work =
 - Histogram (paw/root): 2GHz + 1GB + 200GB
 - Code develop/debug: add CDF sw (AFS, 2GB cache)
 - Job submission: GRID/Kerberos
- Interactive work locally means one of the two:
 - Heavy Desktop (2CPU, 4 disks, frequent upgrade)
 - Light Desktop (~X-term , more flexible)
 - ☞ mini-farm with non-batch access

BaBar needs

- Continued support for (growing) farms in Padova and Roma
 - Mostly ran by experiment personnel (both physicists and computing)
 - BaBar people take care remotely of farm needs
 - ☞ decentralized manpower (make everybody usefull)
 - ☞ centralized hardware (make workload minimal)
 - ☞ single facility (make any work serve everybody)
- No other farms planned
- Support for interactive work everywhere
 - My guess is that this is similar to CDF
- GRID soon

Esperimento in Running: l'esperienza di CDF

- **Stabilita'**

- Ambiente definito, costante, funzionante

- ☞ possibile pianificare lavoro in base a cosa c'e'/ci sara'

- ☞ niente promesse a vuoto please

- **Continuita'**

- Non si possono perdere giorni od ore di lavoro (sul rivelatore e' sempre un'emergenza) perche' si e' rotta una macchina/disco o perche' si cerca di fare meglio quando bene era sufficiente

- **Minimo garantito**

- Servizi essenziali innanzi tutto

- Nuove soluzioni: con cautela

- Sicurezza: non gabbia

Il "minimo garantito ovunque per tutti"

- Net (also wireless)
 - Security, but also access to and from outside
 - ☞ data copy, GRID
- E-mail 24x7 from anywhere
 - IMAP, IMHO, backup
- Windows+Linux, ps+pdf
 - Calcolo : Linux
 - Talks+Ammin: PPT+Excel
 - Docs: Latex + Word
- Desktops
- Home directories
- Scratch areas
- Web server
- Mail list server
- Local clusters must be robust
 - Dual hosts, RAID, Backup, UPS
- Printers
- Laptop
- Eterno problema: il manpower
- Copiare ? Centralizzare/clonare (e-mail, print, web, MS-Office) ?
 - Eg: my web page = directory AFS, PPT = Citrix @ FNAL

CDF wishes: software

- Large, neverending need for software development
 - C++ - Java - Html - Perl - Python - Root - SMQL - cvs - gdb - (x)emacs - Netscape - Memory checker - code optimisation...
- Aiuto = Consulenza
 - esperti dal servizio calcolo ?
- Aiuto = Sviluppo
 - possibile passare al servizio calcolo lo sviluppo di tools, GUI's, pezzetti di package ?
 - non tanto un bisogno quanto un'opportunità
 - il calcolo come la meccanica e l'elettronica: fornire parti critiche e qualificanti di un esperimento
 - Succede in realta', ma in modo un po' personale, scoordinato e con grandi differenze di qualita'

Conclusioni

- "The computer is the network"
- Architettura:
 - Most computing work done in a few dedicated places
 - ☞ Alcuni servizi di calcolo (Padova e.g.) hanno le competenze per rendere possibili farms "di produzione", ma..
 - ☞ Farm = heavy work → minimize, optimize
 - Sezione = portale ai servizi
 - ☞ Mail - web - GRID
 - ☞ Unix + Windows
- Software:
 - Gli esperimenti si arrangiano grazie ai laboratori (esteri)
 - ☞ molta frammentazione, nessuna soluzione diventa di tutti
 - C'e' spazio nell'INFN per un "servizio di consulenza e sviluppo sw di alto livello"

Spare Slides

Da qui in poi sono solo slides spares,
miscellanea, etc.

CDF numbers (2004) and analysis steps

- Data: 500TB/year
- Reco: 5sec/ev
- Ana: 1sec/ev


Step	INPUT	OUT vs IN	Who When	Where	What hardware
Reconstruction	200TB	x2	CDF (2~3 times)	Farm at FNAL (now)	200 duals
Skim	400TB	1/4	~10 Groups 1/week	Same as above	Same as above
Batch Analysis	1TB	1/10 ~ 1/100	~200 Users 1/day	Farm at FNAL (to do)	O(500) duals +O(200)TB
Interactive Analysis	10GB ~ 100GB	1/1000 plots !	~200 Users 1/hour	Home	1 dual +200GB / user

- Everything grows x2/1.5year to keep up with data (Moore)
 - Mostly replace hw every 3 years, ~same number of boxes

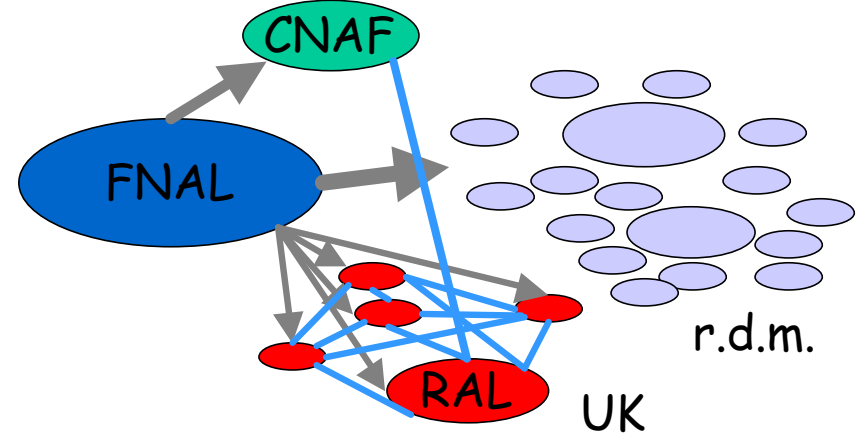
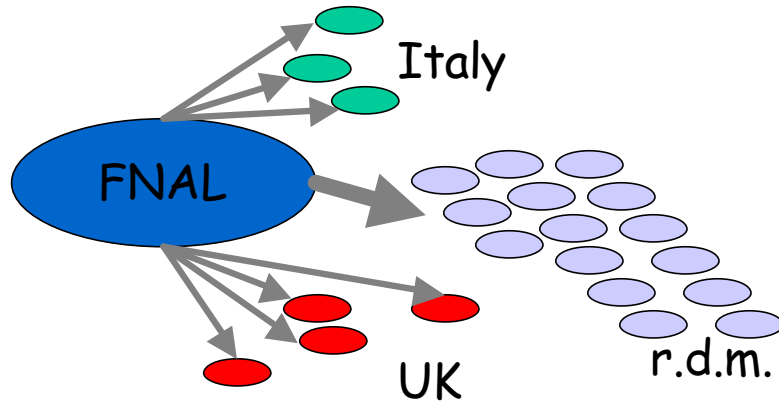
Computing model for Italian CDF data analysis

- Until CNAF Regional Center is up:
 - Everything $O(10\text{TB})$ or more at FNAL
 - Make n-tuple at FNAL and copy to Italy: $O(100\text{GB})/\text{user}$
 - Limited copy of frequently used data sets: $O(1\text{TB})/\text{site}$
 - Hardware in Italy INFN sites:
 - ☞ access to FNAL
 - ☞ interactive work
 - ☞ mini-farms for local data sets: $O(10\text{CPU})$
- After
 - Everything $O(100\text{TB})$ or more at FNAL
 - Copy of analysis data sets to CNAF: $O(10\text{TB})$
 - Hardware in Italy INFN sites
 - ☞ as above but NO Mini-farms
 - ☞ maybe also very limited interactive (4~5 times less)

Computing Model evolution



CDF



BaBar

