Computing Issues

(from a BaBar Perspective)

Super B Factory Workshop

Honolulu, Hawaii

21 January 2004

Rainer Bartoldus

SLAC

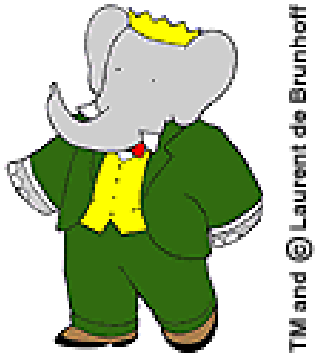

Outline

- **What this talk is about**
  - Computing needs for a future B factory delivering luminosities between $2 \cdot 10^{35}$ $cm^{-2}s^{-1}$ and 10^{36} $cm^{-2}s^{-1}$
  - Look at extrapolation/scaling of current ($\sim 10^{34}$ $cm^{-2}s^{-1}$) resources
  - Expected issues for Event Filtering, Reconstruction, Simulation
  - Some thoughts on Data Distribution and Analysis
- **What I will not discuss (but feel I should)**
  - Issues for Detector electronics/readout, DAQ hard-/software
  - Implications for a viable Trigger
    - (cannot be discussed without a clear definition of Physics goals, priorities, requirements on efficiency and knowlegde thereof etc.)

Y(4S) Cross Sections

- **Production cross sections**

Bb	1.05 Nb
Cc	1.3 Nb
Ss	0.35 Nb
Dd	0.35 Nb
Uu	1.39 Nb
Tt	0.96 Nb
Mm	1.16 Nb
Ee	50 Nb

Data Processing

- **Two Pass System**
  - First pass on a (small) fraction of the events performs **Prompt Calibration**
    - Constants are loaded into Conditions Database to be applied in the second step
  - Second pass on all events performs **Event Reconstruction**
    - Tracking, Clustering, Particle ID etc.

Prompt Calibration

- Calibration essentially tracting variations in time (alignment, beam spot etc.)
- Involves preliminary reconstruction of a fraction of the data
- Currently using a fixed rate of 1 Hz each of Bhabha, radiative Bhabha, mu-pair, hadronic events
- One farm (32 nodes) capable of 600 pb^{-1}/day
- Does not need to scale with luminosity
- "Rolling" calibrations (where the results of one run feed into the next) require this step to be **sequential**
  - Bootstrap at the beginning of a running period
- Could move this into the online

Event Reconstruction

- Uses input from previous step
- Highly **parallelizable**
- One farm (32 nodes) typically capable of 150 pb^{-1}/day
- Current model is for capacity to tract peak luminosities
- By allowing for latency, could relax this to tract integrated luminosity instead
- Can be repeated with improved reconstruction code / calibrations / understanding
  - In a regime where accumulated data sample tends to double in time, reprocessing remains at comparable effort
  - When luminosity curve is flat, reprocessing requires extra (increasing) resources

Simulation Production

- **Current Generator Mix (BaBar)**
  - 2 billion events total currently
  - Generic *BB (3:1)* (30%)
  - Continuum *(1:1)* (35%)
  - Signal *B* (15 %)
  - Other $(\tau^+\tau^-, \ ...)$
  - Change in the composition is an option
    - Verify 3:1 data/MC ratio for generics
    - Fraction of Continuum MC

SP Resources

- **Production Resources**
  - Distributed over >25 sites
    - ¼ at Tier A centers, rest at universities (Tier C)
      - At Tier A, running parasitically on free CPU cycles (where CPU is dedicated as opposed to batch)
  - Presently around 2000 CPUs
    - Capacity of 60 M events/week
  - Not costed as mostly universities
  - For higher luminosities assume no cost as well?

Current Resources

- **BaBar's Five Tier A Centers**
  - **SLAC**
    - ~3000 CPUs, 200 TB disk
  - **IN2P3** (Lyon, France)
    - 350 CPUs, 40 TB disk (plus dynamic staging)
  - **RAL** (UK)
    - 300 CPUs, 50 TB disk
  - **INFN** (Padova+Bologna, Italy)
    - 350 CPUs, 25 TB
  - **GridKa** (Karlsruhe, Germany)
    - 150 CPUs, 30 TB disk

Cost Extrapolations

- **Based on BaBar's current estimates for 2008**

$M	3.3x10^34	2x10^35	10^36
Tape Media	0.7	1.7	6.5
Prod. Disk	0.2	1.8	9.3
Prod. CPU	0.3	3.6	19.4
Add. Tape Drives	*2*	*49*	*275*
Total Prod.	0.5	5.4	28.7
Anal. Disk	0.3	2.2	11.3
Anal. CPU	1.2	5.3	25.2
Add. Tape Drives	*9*	*52*	*261*
Total Anal.	1.5	7.5	36.5

  - Assuming 7 nb processed cross section (16 nb logged)
  - Moore's Law of x1.6/year (disk and CPUs)
  - Has to be taken with a pinch of salt

Distributed Computing

- CPU and disk resources can be expected to follow Moore's Law
- One thing people expect we will have in abundance in the future is network bandwidth
  - Should take advantage of that in exploiting available resources
- BaBar's example shows that the distributed model works
  - Opportunity to reach out to and involve more funding agencies
- But we will not be alone in the world
  - Right now benefiting from shared resources built for LHC and others
  - In the future, will have to compete for shares with Atlas, CMS, ...
- GRID should help to tap a larger pool of resources

Summary and Initial Conclusions

- **Computing in a Super B-Factory**
  - Relying on Moore's Law, but not enough
  - Pushing the start date may gain a factor 1.5
- **$2 \cdot 10^{35}$ cm^{-2}s^{-1}**
  - This may be achievable in the anticipated time frame if Physics goals are well defined
- **10^{36} cm^{-2}s^{-1}**
  - This is likely to require a rather different approach from how we are doing things today