SuperB Computing Resource Estimates (v2)
a.k.a.
Kaffeesatzlesen
“Reading the Coffee Sediments”

SuperB Distributed Storage Meeting 15.2.2012

S. Luitz

An Attempt to Estimate the SuperB Detector Computing Needs

- Based on extrapolation of the BaBar computing model
  - Work by R. Mount, F. Bianci, M. Morandin, S. Luitz
  - Limited Scope
  - Starting 2016 (1st year of data taking), first large purchase in 2015
  - Does not include computing for TDR, detector design, data challenges, engineering runs, accelerator, theory, etc.
- Limited Precision due to many assumptions about
  - Overall computing model (i.e. BaBar-like)
    - Raw, mini, micro, skimming, tape
    - # of copies, fraction of data on disk, amount of Monte Carlo, etc.
  - Raw and reconstructed event sizes
    - i.e. amounts of data per unit
    - integrated lumi CPU per lumi
- Cost estimate – even more uncertainties
  - If exponential growth in computing performance continues, the starting year has a big impact
  - How well can we take advantage of new CPU and storage architectures / paradigms?
- **But: Useful for setting the scale of the adventure!**

Basic Assumptions

- Data taking starts in 2016
- Raw Event size ~200kByte (= 6 x BaBar)
- Mini/Micro event size = 2 x BaBar
- CPU / unit lumi:
  - 4 x BaBar for reco and MC (reflecting larger raw data size)
  - 3 x BaBar for everything else
  - Hand-waving arguments pro & con
    - Data size, combinatorics, etc.
    - Exact value has impact on cost, just making a guess here
- 2 copies of raw data
- Skim expansion factor: 5
- Some fraction of Mini on disk (100% → 10%)
- Equivalent amount of MC "lumi"
- Raw data stored on tape
  - Or media with tape-like cost
- No wide-area networking in cost and resource planning included

BaBar Baseline and Scaling

Year	BaBar	2015	2016	2017	2018	2019	2020	2021	2022
Size of Raw Data(TB/ab-1)	875	6	6	6	6	6	6	6	6
Micro Data Size(TB/ab-1)	42	2	2	2	2	2	2	2	2
Mini Data Size(TB/ab-1)	80	2	2	2	2	2	2	2	2
Micro MonteCarlo Size(TB/ab-1)	51	2	2	2	2	2	2	2	2
Mini MonteCarlo Size(TB/ab-1)	78	2	2	2	2	2	2	2	2
CPU: Phys. Analysis of data (KSpecInt2000/ab-1)	1200	3	3	3	3	3	3	3	3
CPU: Phys. Analysis of MonteCarlo (KSpecInt2000/ab-1)	1300	3	3	3	3	3	3	3	3
CPU: Data Reconstruction (KSpecInt2000/10^36 cm^2sec^-1)	22100	4	4	4	4	4	4	4	4
CPU: MonteCarlo Production (KSpecInt2000/10^36 cm^2sec^-1)	70000	4	4	4	4	4	4	4	4
CPU: Skimming of Data (KSpecInt2000/10^36 cm^2sec^-1)	10200	3	3	3	3	3	3	3	3
CPU: Skimming of MonteCarlo(KSpecInt2000/10^36 cm^2sec^-1)	10000	3	3	3	3	3	3	3	3
CPU: Phys. Analysis of data (KHepSpec/ab-1)	4.8	3	3	3	3	3	3	3	3
CPU: Phys. Analysis of MonteCarlo (KHepSpec/ab-1)	5.2	3	3	3	3	3	3	3	3
CPU: Data Reconstruction (KHepSpec/10^36 cm^2sec^-1)	88.4	4	4	4	4	4	4	4	4
CPU: MonteCarlo Production (KHepSpec/10^36 cm^2sec^-1)	280	4	4	4	4	4	4	4	4
CPU: Skimming of Data (KHepSpec/10^36 cm^2sec^-1)	40.8	3	3	3	3	3	3	3	3
CPU: Skimming of MonteCarlo(KHepSpec/10^36 cm^2sec^-1)	40	3	3	3	3	3	3	3	3

Storage

Luminosity (ist.:pb-1/s, int.: ab-1)								
Peak	0	0.25	0.70	1.00	1.00	1.00	1.00	1.00
Per year	0.00	3.75	10.51	15.01	15.01	15.01	15.01	15.01
Integrated	0.00	3.75	14.26	29.26	44.27	59.28	74.29	89.30
Storage (TB)		2016	2017	2018	2019	2020	2021	2022
Raw Data (integrated)	0	39,395	149,702	307,282	464,863	622,444	780,024	937,605
Micro Data (integrated)	0	1,261	4,790	9,833	14,876	19,918	24,961	30,003
MiniData (integrated)	0	600	2,281	4,682	7,084	9,485	11,886	14,287
Micro MC (integrated)	0	1,531	5,817	11,940	18,063	24,186	30,310	36,433
Mini MC (integrated)	0	585	2,224	4,565	6,907	9,248	11,589	13,930
Skims (Data and MC, 1 year))	0	3,489	9,770	13,957	13,957	13,957	13,957	13,957
User Data	0	356	1,352	2,774	4,197	5,620	7,043	8,465
Production Buffer	0	377	1,055	1,507	1,507	1,507	1,507	1,507
One reprocessing cycle of micro&mini (data and MC)	0	0	1,883	7,157	14,691	22,225	29,759	37,293
A reskimming cycle (data and MC)	0	0	3,489	13,259	27,216	41,174	55,131	69,088
Event Store Buffer (50% of n-2 reprocessing cycle)	0	0	0	0	2,686	10,208	20,954	31,699
Storage Total	**0**	**47,594**	**182,363**	**376,958**	**576,047**	**779,971**	**987,119**	**1,194,267**
Disk Storage (net)	0	8,199	27,289	44,635	59,595	67,062	80,124	93,187
Disk Storage Total (including technical reserve)	**0**	**9,428**	**31,382**	**51,331**	**68,534**	**77,121**	**92,143**	**107,165**
Tape Storage Total	**0**	**43,728**	**168,049**	**348,235**	**530,680**	**713,125**	**895,571**	**1,078,016**

Storage grows from O(50PB) to O(1000PB) over 5 years

CPU

Luminosity (ist.:pb-1/s, int.: ab-1)								
Peak	0	0.25	0.70	1.00	1.00	1.00	1.00	1.00
Per year	0.00	3.75	10.51	15.01	15.01	15.01	15.01	15.01
Integrated	0.00	3.75	14.26	29.26	44.27	59.28	74.29	89.30

CPU (kHEPSpec)		**2016**	**2017**	**2018**	**2019**	**2020**	**2021**	**2022**
Physics Analysis of data	0	54	205	421	638	854	1,070	1,286
Physics Analysis of MC	0	59	222	457	691	925	1,159	1,393
Beam data reconstruction	0	88	248	354	354	354	354	354
MonteCarlo generation and processing	0	280	784	1,120	1,120	1,120	1,120	1,120
Skimming of Data	0	31	86	122	122	122	122	122
Skimming of MC	0	30	84	120	120	120	120	120
Reprocessing of beam data (previous years)	0	0	88	336	690	1,043	1,397	1,750
Regeneration of MC (Previous years)	0	0	280	1,064	2,184	3,304	4,424	5,544
Reskimming of reprocessed data	0	0	46	174	358	542	725	909
Reskimming of reprocessed MC	0	0	45	171	351	531	711	891
CPU Total	**0**	**542**	**2,088**	**4,339**	**6,627**	**8,914**	**11,202**	**13,489**

Equipment Purchases

- Assuming doubling of performance/capacity per unit cost every two years
  - "Moore's law"
  - Need to be able to take advantage of CPU architectures
- Hardware lifecycle assumptions
  - Typical IT lifecycles today are ~4 years
  - Environmental concerns and flattening of Moore's law may push this to 5 (We've already seen an increase from 3 to 4 years)
  - For this exercise: Assume buy in year 3 so that old equipment can be shut down in year 4
- No copying of old tapes, no tape infrastructure cost included

Industry Predictions

Year	**2011**	**2012**	**2013**	**2014**	**2015**	**2016**	**2017**	**2018**	**2019**	**2020**	**2021**	**2022**
Cost of a tape Silo (kEuro)	350.00	350.00	350.00	350.00	350.00	350.00	350.00	350.00	350.00	350.00	350.00	350.00
Cost of a tape drive (kEuro)	20.00	20.00	20.00	20.00	20.00	20.00	20.00	20.00	20.00	20.00	20.00	20.00
Cost of a tape (kEuro)	0.08	0.08	0.08	0.08	0.08	0.08	0.08	0.08	0.08	0.08	0.08	0.08
Tape size (TB)	1.00	2.00	2.00	5.00	5.00	10.00	10.00	20.00	20.00	40.00	40.00	80.00
Tapes per Silo	6000.00	6000.00	6000.00	6000.00	6000.00	6000.00	6000.00	6000.00	6000.00	6000.00	6000.00	6000.00
Tapes/tape drive	200.00	200.00	200.00	200.00	200.00	200.00	200.00	200.00	200.00	200.00	200.00	200.00
Cost of disk (kEuro/TB)	0.60	0.42	0.29	0.21	0.14	0.10	0.07	0.05	0.03	0.02	0.02	0.01
Cost of CPU (kEuro/KSpecInt2K)	0.06	0.04	0.03	0.02	0.01	0.01	0.01	0.00	0.00	0.00	0.00	0.00
Cost of CPU (kEuro/KHepSpec)	15.00	10.50	7.35	5.15	3.60	2.52	1.76	1.24	0.86	0.61	0.42	0.30

Cost

Purchase (without replacement) [kEuro] | Year	2015	2016	2017	2018	2019	2020	2021
Disk	1,358	2,214	1,408	850	297	364	255
Tape	700	995	1,441	730	730	365	365
CPU	1,950	3,899	3,973	2,826	1,978	1,385	969
Replacement (4 years)							
Disk [kEuro]	0	0	0	466	759	483	292
Tape Media [kEuro]	0	0	0	0	0	0	0
CPU [kEuro]	0	0	0	669	1,337	1,363	969
Total [kEuro]	**4,008**	**7,108**	**6,822**	**5,540**	**5,102**	**3,959**	**2,849**

- Total computing hardware/media cost for SuperB in this model: between 5-10MEuro / year
- Moore's Law seems to balance luminosity growth

Conclusion

- I'm getting really tired of talking about extrapolations from BaBar
  - But STILL doing it
  - We still do the same for Online, too ☺
- For the TDR we need to work on reducing the uncertainties for SuperB
  - Computing model
  - Event sizes, storage and CPU needs