



C M S

Minutes of the 12th RESOURCES REVIEW BOARD Meeting

(Held at CERN on 23rd April 2001)

Present:

Europe:

H. Borns (Bundesministerium für Bildung, Wissenschaft und Kultur, Wien), C. E. Wulz
J. Lemonne (NFWO, Brussels), J. Sacton (FNRS, Brussels);
G. Anchev (Ministry of Education and Science, Sofia);
D. Denegri (Ministry of Science and Technology, Zagreb);
D. O. Riska (Helsinki Institute of Physics, Helsinki), J. Tuominiemi;
J. Feltesse, P. Debu (CEA-Saclay, Gif-sur-Yvette), P. Brossier, J. Rander;
G. Wormser (IN2P3, Paris), B. Ille;
J. Richter, A. Lindner (BMBF, Bonn);
G. Vesztergombi (KFKI-RMKI, Budapest);
M. Calvetti (INFN, Roma), U. Dosseli;
J. Królikowski (State Committee for Scientific Research, Warsaw);
F. Bello (JNICT, Lisboa), G. Barreira;
F.E. Grishaev, V.I. Savrin (Ministry of Science and Technology, Moscow);
A. N. Sissakian (Dubna), I. A. Goloutvine;
A. Ferrer (CICYT, Madrid), M Cerrada;
R.A. Eichler (PSI, Villingen), F. Pauss (ETHZ);
I. F. Corbett (PPARC, Swindon), R.M. Brown.

N.America:

M. Goldberg (NSF), C. Laugguth;
J.R. O'Fallon (DOE), D. Green, J. Yeck.

Asia:

S. Bhave (Department of Atomic Energy, Mumbai), S. N. Ganguli;
C-T. Hu, S.Hou (NSC, Taipei).

CERN:

R.J. Cashmore (chairman), E.M. Rimmer (secretary),
L. Maiani, M. Delfino, G.V. Goggi, H. F. Hoffmann, C. Jarlskog, J. May,
A.J. Naudi,
E. Tsesmelis, E. van Hove.

CMS:

M. Della Negra, A. Ball, D. Blechschmidt, L. Foà, A. Hervé, D. Stickland, M.
Pimiä, H. Rykaczewski, T. Virdee.

Apologies:

University of Cyprus, Nicosia
Pakistan Atomic Energy Commission, Islamabad

Plenary Session

1. Welcome (*R.J. Cashmore, Director for Collider Programmes*)

The Chairman, R.J. Cashmore, welcomed RRB delegates of ALICE, ATLAS, CMS and LHCb. He announced that, having failed to take place at the last meeting, visits to ATLAS and CMS areas are scheduled during the present meeting. He also asked delegates to fill out a questionnaire concerning the use of e-mail and the Web for LHC RRB business.

2. Status of the LHC (*L. Maiani, Director General*)

Professor Maiani started with an artist's impression of the LHC machine *in situ*. Reporting on progress in underground/overground civil engineering, he highlighted the impressive work at the ATLAS and CMS areas. Excavation and lining of the ATLAS vault is nearing completion, after which excavation will begin of the cavern itself, a 6-storey-deep hole below the vault. At CMS, a huge reinforced central pillar is being constructed to support the roof before excavation of the two adjacent caverns can start.

The DG showed a picture and training curve of the first superconducting dipole ready for installation, one of pre-series of 30 dipoles per firm ordered last year. The magnet experienced its first quench just below the nominal field of 8.3 Tesla, reaching the ultimate field of 9 Tesla without a further quench. The DG expressed the wish that the remaining 1235 dipoles should be equally well behaved.

Following discussions between those responsible for the experiments and for the machine, the LHC commissioning schedule has been brought into line with the current status:

April – August 2004	first octant test
March 2005	last dipole delivered
end 2005	ring closed and cold
February 2006	first circulating beams
April 2006	1 month pilot run with collisions
May – July 2006	shutdown
August 2006 – February 2007	7 months of p-p collisions at $L > 2 \times 10^{33}$
April 2007	6 weeks of Pb-Pb collisions

The DG remarked that the revised schedule is not so different from the one drawn up in 1996 when difficulties such as the adverse geology and pillar excavation in the CMS area were as yet unknown. He also noted that not having winter shutdowns in 2005 and 2006 will affect CERN's electricity contracts.

A realistic cost-to-completion estimate for the machine should be possible by the end of the year, after some big contracts for dipole assembly have been adjudicated in

September. The DG concluded by pointing out that costs for LHC computing and detector Maintenance and Operation, to be discussed at this and at the October RRB meetings, should also be fairly well known by the end of 2001.

3. Maintenance and Operation (*R.J. Cashmore*)

Cashmore opened by reminding delegates of his presentation on M&O at the October 2000 RRB meeting. This had been followed by several bilateral conversations and a paper (*RRB-D 2001-04*) had been distributed in preparation for detailed discussions at the current meeting.

He drew attention to the scale of LHC experiments, which, taken together, are some 4 times bigger than the LEP experiments in capital investment, the number of people involved and M&O costs. During 2000, a CERN Working Group on M&O had collated input from the experiments, from LEP physicists and from CERN experts to prepare preliminary cost estimates for the years up to 2007.

The main principles concerning M&O cost sharing were developed by the CERN Management and presented to the Scientific Policy Committee in March 2001. The SPC expressed its general support for an approach that would: follow the best practices used at LEP, HERA etc.; aim for uniformity and transparency; establish total costs including core manpower; share general costs by established scientific author; leave sub-detector maintenance with the constructors; allow payment in cash and in kind but with a minimum cash contribution; take into account contributions that Member States and some non-Member States have made to building the LHC machine. These basic principles should be discussed during these current RRBs so that a draft MoU for M&O from 2003 onwards can be prepared for October. Interim arrangements for 2002 will have to be agreed in October, as substantial M&O costs are already being incurred.

Cashmore then listed the cost categories; A for common operations, Common Fund items and collaboration-wide services, B for maintenance of items provided by sub-sets of the collaboration, and C for items that are the responsibility of host laboratory. ICFA guidelines, signed MoUs for LHC detector construction and CERN's General Conditions for Experiments had been taken into account when categorising cost items. In essence, CERN must provide safe, bare experimental caverns and ancillary buildings, and the collaborations must cover costs incurred because detectors are installed.

Cashmore emphasised that present cost estimates are not yet reliable for several reasons. All items must be examined for completeness, correctness and categorisation and a clear definition must be given for each one so that costs can be correctly estimated and double counting avoided. The collaborations must bring their cost estimates in line with the 2006 machine start-up schedule and refine them, if necessary, to fit the tighter item definitions. In particular, sub-detector maintenance costs need careful revision. That having been done, scrutiny groups with some members nominated by the RRBs must examine the final estimates to assure the RRBs that they are reasonable.

The preliminary total M&O cost estimates are some 18MCHF in 2002 rising to 65MCHF in 2007. These numbers reduce to 13MCHF and 44MCHF respectively after subtracting category B costs and CERN's C and A costs.

Cashmore reported a widespread preference for sharing A costs by qualified scientific authors of publications or Technical Design Reports, that is, based on exploitation of the detector. If B costs are based on retaining responsibility for sub-systems, cost sharing details will be left to the collaboration. However, these sub-systems are highly complex, and responsibilities and commitments will have to be clearly reported to the RRBs. A common recovery plan should be envisaged in case of a major disaster, an issue that is linked to CERN's Insurance policy, presently under review.

Addressing the fact all Member States and some non-Member States have contributed to building LHC, Cashmore mentioned the idea of rebates, whereby CERN pays a fraction of the A costs of Agencies and Institutes belonging to those States. The SPC had been enthusiastic about this suggestion. The level of the payments and the formulae by which they would be apportioned would require endorsement by the CERN Council.

Cashmore then described the proposed procedure for dealing with M&O costs. In April of year N-1, the RRBs would be given preliminary estimates of A and B costs for year N and a forward look for N+1, N+2 and N+3. They would also receive the final accounts of A costs and final reports of B costs for year N-2. In October of year N-1, the RRBs would be given scrutinised estimates for A and B costs for the years N, N+1, N+2 and N+3. They would approve final allocations for A costs and note final arrangements for B costs for year N.

Cashmore expressed his hope that the current RRBs would discuss arrangements for handling M&O in 2002 and agree that draft M&O MoUs be prepared for examination in October. Final versions of the MoUs could then be presented for approval to the RRBs in April 2002 and thereafter circulated for signature.

In conclusion, Cashmore stated that M&O is an important issue for commissioning and exploiting LHC detectors. Serious costs are already being incurred and it is becoming urgent that CERN, the collaborations and the RRBs establish formal arrangements for handling these costs, to guarantee the success of the LHC experiments.

Discussion

I.F. Corbett (GB) pointed out the considerable overlap between the experiments on issues of both M&O and LHC computing. He asked whether the October meetings could be organised to take account of that and optimise the use of what will be a very limited amount of time. Cashmore agreed, adding that it is important to best use the time before as well as at the October meetings.

4. The LHC Computing Review *(S. Bethke, Steering Committee Chairman, LHC Computing Review)*

Computing Review Steering Group Report (RRB-D 2001-03) is available at <http://lhc-computing-review-public.web.cern.ch>

LHC computing is a vital prerequisite for the success of the experiments and a potential source of spin-offs yet unimagined. Bethke explained that LHC offline computing was not included in detector MoUs because of severe uncertainties in extrapolating more than 5 years ahead in this rapidly developing field. It is now both appropriate and necessary to finalise LHC computing plans, start serious prototyping and secure the resources needed for timely completion.

The LHC Computing Review was conducted by three independent panels reporting to a Steering Committee: a WorldWide Analysis / Computing Model panel (chaired by D. Linglin, CC-IN2P3/CNRS), a Software Project panel (M. Kasemann, FNAL) and a Management & Resources panel (M. Calvetti, INFN Florence). Membership encompassed representatives of the four experiments and CERN IT Division, and experts from around the world.

LHC computing is an unprecedented challenge for the HEP and IT communities, as Bethke illustrated with statistics that included total annual storage requirements of 7 PBs (7×10^{15}) for raw data and 3.2 PBs for simulated data. Each year, Tier0+Tier1+Tier2 centres (see later) will need a total tape storage capacity equivalent to 40 million CD-ROMs and disk storage equivalent to 140 thousand 75 GB disks. Their combined CPU capacity will equal that of 360 thousand of today's PCs with a total Tier0 \leftrightarrow Tier1 WAN bandwidth of 5000 Mbps to serve the four experiments. The Review accepted the scale of these resource requirements estimated by the experiments.

The World Wide Analysis / Computing Model panel recommended adopting the distributed, hierarchical model developed by MONARC, a collaborative effort on Models of Networked Analysis at Regional Centres. This model consists of a Tier0 centre at CERN, to store all raw data and perform some reconstruction, connected to several regional/supranational Tier1 centres for analysis, Monte Carlo generation and the like (one Tier1 centre being at CERN). Each experiment will require the resources of about five Tier1 centres plus a larger number of similar but smaller national/intranational Tier2 centres. The Panel assumed that institutional Tier3 facilities and end-user Tier4 workstations will anyway be available and so did not cost them. GRID technology was recognised as suitable for the transparent and efficient use of these distributed resources, and the need was identified for affordable networking at 1.5 - 3 Gbps per experiment by 2006.

The Software Project Panel called for joint efforts and common projects between the experiments and CERN-IT, with support for widely used packages. A matrix showing who is developing/maintaining/using which software packages illustrated the varying popularity and vulnerability of these products. Data challenges of increasing size and

complexity were seen as essential steps in developing production software and CERN was asked to back the transition to OO programming. Areas of concern were the limited maturity of planning and resource estimates, insufficient development and support of simulation packages, and inadequate support for and evolution of analysis tools.

The Management & Resources panel underlined that current cost estimates are based on evolutions forecast by the PASTA committee, the technology tracking team for processors, memory, storage and architectures set up by IT Division and the LHC Computing Board. This foresees logarithmic increases in capacity/performance and decreases in unit cost with time for computer hardware, trends that are currently observed but not guaranteed to continue. The hardware costs of the initial Tier0+Tier1+Tier2 centres are estimated at 240 MCHF, a third of which is for the CERN-based Tier0+Tier1 pair. These numbers will have to be reviewed every couple of years because of uncertainties in the actual performance of the LHC machine, the detectors, triggers, backgrounds and so on. Assuming that LHC starts up in 2006 and reaches design luminosity in 2007, the investment will have to be equally spread through 2005, 2006 and 2007.

A major concern is the chronic understaffing of the teams producing core software (which is everything except physics codes and GRID middleware). Shortfalls range from 28 FTEs already in 2000, reaching 42.5 FTEs in 2002 and slightly falling to 35.5 in 2005, a problem that must be addressed by the collaborating institutes. In the same context, the Review found the planned reduction of CERN-IT staff to be incompatible with providing CERN-based LHC computing and software support.

Estimates for Maintenance and Operation of the LHC computing system are based on rolling replacement within a constant budget. About 30% of the initial investment would be needed each year, namely some 80 MCHF world wide, a sum which would include the steady evolution of capacity. Bethke noted that a similar approach at LEP had given an increase in computing capacity of a factor of 1,000 during the 15 or so years between 1985 and switch-off.

To develop the final LHC computing system, a common prototype must be set-up as joint project between the experiments and CERN-IT with the participation of some major Tier1 and Tier2 centres. By 2003/4, the prototype should reach about half the complexity (not capacity) of one LHC experiment. The estimated cost of the prototype, about 18 MCHF, is not included in the initial investment costs and so an agreement is urgently needed on how to construct and finance it.

The Review strongly recommended setting up an LHC Software and Computing Steering Committee (SC2) composed of the highest level of computing management in the experiments, CERN-IT and regional centres, to steer the development and deployment of the entire system. The SC2 should establish Technical Assessment Groups (TAGs) to launch specific tasks and projects.

Each collaboration must prepare an MoU for LHC computing that defines the overall required funding and agreed responsibilities. As an interim measure, IMoU's or software agreements should be in place by end of 2001 if possible.

The Review has shown the enormity of the LHC computing challenge and Bethke reiterated the non-negligible chance of spin-offs. It has underlined the crucial importance of proper funding, planned development and timely realisation of the entire LHC offline computing system, and its later maintenance and operation. As Bethke had earlier remarked, without adequate and appropriate computing facilities, the LHC machine and its detectors will be of little use.

Discussion

Asked by S. Bhave (IN) whether cost sharing had been addressed, Bethke replied that the Review had considered costs but not their sharing. Some countries are already setting up Tier1 centres and it will be the RRBs' job to discuss cost distribution as well as how to finance connections to countries with only small national facilities. G. Wormser (FR) remarked that the Review had done a good job defining a common hardware infrastructure. He asked whether, to mitigate the lack of manpower, there is a schedule for choosing between the several software options still open and what role CERN-IT will be able to play in the decision process. Bethke acknowledged the importance of these remarks, adding that software choices had been one of the most difficult issues faced by the Review. An attempt to streamline the use of packages had not had much success because most experiments have already invested considerable development effort based on particular packages. While many packages are used by several experiments, some are used by only one or two and Bethke concluded by saying that common attempts at streamlining had been started and must continue.

CMS Resources Review Board Meeting

1. Introduction (*R.J. Cashmore, Director for Collider Programmes*)

The Chairman, R.J. Cashmore, welcomed delegates and reported that since the last meeting an MoU has been signed with the Funding Agency in Taipei. He reminded delegates to fill out the questionnaire concerning the use of e-mail and the Web for LHC RRB business.

2. Approval of the minutes of the 11th Meeting

The Minutes were **approved** without changes.

3. Status of the experiment (*Spokesperson M. Della Negra*) (*CMS RRB-D 2001-102; CMS RRB-Tr 2001-102*)

News from the collaboration

Since the last RRB, three US groups have joined the Silicon Tracker Project as Full Members of CMS: Kansas State University, the University of Kansas and the University of Santa Barbara. Two new Associate Members are the Beijing Glass Research Institute, China, with an interest in PbWO_4 crystals, and the Department of Computer Science, University of Strathclyde, UK, interested in software verification, validation and process improvement.

Iran has applied for Full Membership: 6 Professors and 3 Ph.D. students from the Institute for Studies in Theoretical Physics & Mathematics, Tehran, are interested to work on the Forward Calorimeter (HF). A CMS Workshop will be held in Iran at the end of May. The DG has made an official visit and a Cooperation Agreement and Protocol between CERN and Iran is in preparation. The Florida Institute of Technology has also applied for Full Membership to work on the Forward Calorimeter (Mohammad Mohammadi Baarmand, one of the Iranian Professors, is presently there.)

Discussions continue with two other new countries, Brazil and Ireland.

Status of the sub-systems

Della Negra highlighted two key dates in the updated LHC schedule. In April 2006 a 1-month pilot run will provide first proton collisions at a luminosity of $5 \times 10^{32} - 2 \times 10^{33}$. In August 2006 a 7-months physics run will begin at a luminosity $\geq 2 \times 10^{33}$ giving an integrated luminosity $\geq 10 \text{ fb}^{-1}$ for ATLAS and CMS combined, enough to discover the Higgs boson anywhere between the LEP limit of 115 GeV and 1000 GeV. The question is how much of the CMS detector will be ready.

Civil engineering

An extra delay of 4 months was announced in January 2001 because of problems with the supporting pillar between the underground experimental cavern UX5 and service cavern US5. This brings the total delay to 9 months with respect to the contractual planning and means that UX5 and US5 will now only be delivered in April 2004. This has serious implications for installation and commissioning (see the later presentation of A. Ball) and CMS is including contingency for even further delays.

Some UX activities can be moved to the surface, albeit at the cost of extra resources. However, activities foreseen for the underground service cavern cannot be moved to the surface. The detailed consequences are still being worked out and final milestones will be given to the RRB in October.

Magnet

The magnet is on schedule for an estimated total cost that is practically unchanged (122.1 MCHF). Major contracts for some 92 MCHF (~ 75%) have been placed.

At Point 5, three of the five barrel yoke rings are assembled and the fourth is under construction. Endcap disk assembly started in April. A trial assembly of the first disk YE-1 had been done last November in Kawasaki, Japan before delivery to CERN. For a complete end-cap, 2 such disks with a nose and another 4 disks are needed.

Progress with the super-conducting cable is good. The total requirement is for 21 lengths of 2.65km (5 sections per coil, 4 lengths per section + 1 spare length). To date, 50% of the super-conducting strands (3 x 320km), 4 lengths of Rutherford cable and 3 lengths of real insert have been produced. Electron beam welding reinforcement has been tested with dummy insert and the first full length of final conductor expected within a month.

Magnet testing on the surface should finish by July 2004, slightly later than initially planned but in the shadow of the civil engineering delay.

Tracker

The Tracker Addendum to the MoU has been signed by 21 of the 31 funding agencies. An Engineering Design Review was passed in November 2000. Pre-production of 200 detectors to test mass production is going well. Della Negra showed a picture of a Tracker Outer Barrel 12cmx6cm Si detector fitted with the final rad-hard APV chip. Producing such detectors from big 6" Si wafers was the breakthrough that allowed the MSGCs to be replaced.

Major milestones in 2001 are the system test of final electronics chain and the start of mass production of modules. The tracker should be ready in December 2004, with enough contingency to ensure installation and commissioning in March through July

2005.

ECAL

Della Negra recalled that the motivation for building a crystal calorimeter is to detect low-mass Higgs bosons. A simulation of the two-photon decay of a 100 GeV Higgs (now below the LEP 115GeV limit) shows a clean two-product signature in the ECAL. Coupled with the absence of corresponding signals in the Si tracker, this is convincing evidence for the existence of the Higgs.

The total crystal requirement for the ECAL is 60,000 for the barrel and 16,000 for the endcaps. 6,000 crystals have been delivered on schedule from Russia at a price of 1.6 \$/cc and a contract has been signed for the manufacture of a further 30,000.

Large diameter ingots can now be grown in Russia, allowing two crystals to be cut from one ingot, both for barrel crystals and the larger endcap crystals. This breakthrough has considerably increased the crystal yield per oven. Replies to tender for the remaining 40,000 crystals were opened on April 5. A final contract is being negotiated for an offer that matches the CMS schedule.

The cost of crystals is an on-going concern. More investment has been needed than foreseen to keep to schedule, and the crystals are priced in \$/cc, the prevailing \$/CHF ratio being higher than at the time of the MoU.

HCAL

Both the barrel absorber (HB) and endcap absorber (HE) for the HCAL are on schedule. One half of the barrel has been delivered to CERN from Spain and one endcap has been manufactured and assembled by a Russia/Belarus team.

The forward calorimeter (HF) design has been simplified and it is hoped that the mechanical absorber can be built in Russia. The HF schedule is tight. It will be the first HCAL component to be installed and assembly must be finished by end 2003.

A further concern is a cross-talk problem with the HPD photodetectors that has yet to be understood.

Muon system

Three detector technologies are used in the muon system.

The endcaps are made of Cathode Strip Chambers, mainly built in the US but with one site in China and one in Russia. To date, 30 chambers have been manufactured on schedule in FNAL and 20% of the CSCs should be produced by end 2001. Panel production is progressing in Dubna for ME1/1.

The barrel Drift Tubes are being produced at 4 sites in Italy, Spain and Germany. Two MB2 chambers have been assembled in CIEMAT, Spain, using the final tools and procedures. An important milestone should be passed in May 2001 with at least one DT assembled in Aachen and one in Legnaro (Padova). By the end of 2001 some 20 DTs should be assembled out of the few 100 needed. DT production is running late by about a year, and so a second line will be set up at the site in Torino and the situation will be reviewed at the end of 2001.

Ten pre-production Resistive Plate Chambers for the muon trigger will be assembled by June 2001. Quality checks and control measurements will be made on critical parameters and the 'oil or no oil' question will be addressed.

TriDAS

The Level-1 Trigger TDR, submitted at end of 2000, was recommended for approval by the LHCC in March 2001. Submission of the DAQ TDR has been postponed by one year to November 2002 following the updated LHC schedule.

Two DAQ demonstrators are available. A 32x32 channel Event Builder has been assembled using 64 Intel-PCs interconnected by advanced technology networks, a 64-port Gbit Ethernet switch and a 128-port Myrinet switch. Results are promising and with good indications that the 1000x1000 final system can be built. A complete DAQ column, with Front-End-Driver, Read-Out Unit, Builder Unit, and Filter Unit exhibits no particular problems. CMS plans a progressive implementation of the scalable DAQ system between 2004 and 2008, a timescale with funding implications.

Software and Computing

Software and computing have been completely reorganised into a coherent project called CPT comprising 3 activities: CCS Computing and Core Software, PRS Physics Reconstruction and Selection, and the software and computing aspects of TriDAS.

The transition to OO and C++ code is well underway and functional prototype packages were released in October 2000 for the framework CARF, reconstruction ORCA, and simulation OSCAR based on GEANT4. Millions of events have been reconstructed in ORCA and stored in an Objectivity database for PRS studies. In 2001/2002, >10 million events will be simulated, reconstructed and stored at various computing centres to allow online high level triggers to be demonstrated for the DAQ TDR. Work will continue to make the core framework and related software less dependent on choice of database, the final choice of ODBMS having been moved by one year to December 2002.

Hardware facilities in Tier0, Tier1 and Tier2 prototype centres must now start ramping up to allow testing of distributed production software. Della Negra pointed to the need for R&D in GRID computing. Like the other LHC experiments, CMS regards the lack of professional software engineers to be a very serious problem needing urgent solution.

Milestones, Costs & Funding

With a timetable still based on being ready in October 2005, 89% of the Level1/Level2 milestones are complete to date. This would indicate that the complete CMS detector can be ready for the physics run starting August 2006 (except for the 4th Muon Endcap ME4, staged from the beginning.)

Any limitations come essentially from funding shortfalls or cash flow problems.

	Total cost	Income	Shortfall
In MoU	466 MCHF (full tracker)	457 MCHF (assumed)	9MCHF
LHCC Cost Review September 2000	464 MCHF (Si tracker)	444 MCHF (assured)	20MCHF

The currently assured income of 444 MCHF includes a pessimistic estimate of available US contingency funds for the tracker and a shortfall of 10 MCHF in the Russian commitment. Extra funds of 4.3MCHF have been found (3.1M from the US) to reduce the 20 MCHF shortfall to ≈ 15.7 MCHF.

Additional problems have arisen from: adverse dollar exchange rate variations, extra work surface needed because of the additional 4 month delay in delivery of the underground caverns, some items initially overlooked, e.g. neutron shielding, and some real cost overruns, especially for the ECAL. The total funding shortfall is now:

Lack of funding (in MoU)	15.7
Dollar/CHF exchange rate penalty	10
Extra work on surface	4
Overlooked/new items/cost overruns	6.6
ECAL cost overrun	8
TOTAL	44.3MCHF

This is <10% of the total cost, $\approx 3\%$ due to lack of funding, $\approx 6\%$ cost increase/addition/overrun. Orders are being placed for ECAL crystals and for the tracker, and the installation plan is better defined, so improved cost estimates will be available for the October RRB.

The CMS contingency plan gives priority to critical path items by shifting allocated funds around (within same detector) and finding new funds. A missing 10.5 MCHF have already been covered in this way, but from now on help is needed from the RRB and FAs to complete CMS.

Critical unfunded items at present are Bakelite and electronics for Endcap RPCs. To keep the muon system on schedule Della Negra asked the RRB to approve the borrowing of 0.8 MCHF from the Common Fund to purchase these items. He hoped that, with the help of the FAs, an overall solution to the funding problem could be found within a year. The deliverable items to be covered and the major contributors are:

<i>Deliverables</i>	<i>Cost</i>	<i>Major Contributors</i>
---------------------	-------------	---------------------------

Infrastructure	2.5	CERN, (+)
Surface Work	4.0	CERN, (+)
Tracker	7.4	Italy, CERN, Germany, France, USA
ECAL	18.0	Switzerland, CERN, France, USA
HCAL(HF)	2.0	USA, Russia
Muon DT	2.8	Italy, Germany, Spain
Muon CSC	4.3	USA, Russia
Muon RPC	3.3	Italy, Pakistan, Korea, (+)
TOTAL	44.3MCHF	

To complete the picture, >14 MCHF will be required to upgrade CMS for high luminosity running in 2007/8, essentially to install the staged 4th muon endcap station.

Della Negra noted that extra resources are also needed to cover M&O and Computing, and that first estimates will be discussed in later talks. It appears timely that the experiments and FAs should begin to consider these issues.

Discussion

Asked by S.N. Ganguli (IN) why the funding shortfall has risen so quickly, Della Negra said that the level reported to the RRB for the last 3 years has been ≈ 25 MCHF. New and missing items have added 19MCHF, partly in reaction to the new schedule. Cashmore reminded the RRB that it had asked CMS to estimate the cost to completion, as foreseen in MoU, so as to get good view of the problems.

M. Calvetti (IT) asked how can further cost increases be prevented and what is the strategy to cover extra costs. Della Negra pointed out that CMS funds are now some 50% committed and an even better global view will be possible by October. Remaining uncertainties concern the tracker (with only 2.5M committed out of 77M), the ECAL (still negotiating final crystal cost) and TriDAS and computing where estimates are based on extrapolations. The strategy to cover extra costs is to report to RRB and to go back to those FAs that are providing deliverables to discuss their commitments.

Calvetti queried the sudden urgency of the Bakelite purchase, as the need has been known for a long time. Della Negra explained that the EDR for muon endcap RPCs is in September 2001 and, given a green light, production must start immediately thereafter. It has always been known that the Bakelite was not covered, and nobody is volunteering to pay now. Cashmore remarked that CMS have always had a known shortfall. They have found new collaborators, and continue to do, and have solved many problems inside allocated funds. Now they have reached the limit and need to borrow against a future solution to the funding problem, a situation not untypical of such a project at this stage. Della Negra mentioned that there are 3.6 MCHF in the CF for software, a lower priority item than Bakelite. This is being spent at a rate of 0.4 M per year for licences and maintenance so the 0.8 M could be seen as being taken from this allocation. J. Feltesse (FR - CEA) asked how the software would be affected if no extra money were found. Della Negra replied that effects would only start to be felt in 2005/6, by which time 0.8 M may appear insignificant alongside the problems of computing and M&O costs. He said the real problem is not to borrow 0.8 M but to find 44.3 M.

Cashmore asked whether the RRB would accept the borrowing as legitimate. There being no dissension, it was taken that the **RRB accepted** that 0.8 CHF can be borrowed from the CMS Common Fund for the purposes presented.

Concerning missing funds needed for completion, Cashmore reiterated that the RRB had asked for and received estimates. CMS must now talk to the FAs and present a progress report to the October RRB. The LHCC CMS Comprehensive Review will be also be in October and reported to the RRB. The aim should be to have a final plan ready for April 2002 and the overall strategy should include looking for new collaborators, for extra money and possibly staging the detector.

G. Wormser (FR - IN2P3) observed that it is better to go to governments once only for more money and April 2002 may be too early. Cashmore replied that CMS (and ATLAS) cannot wait forever and delays will mean more requests for borrowing. Maybe April 2002 will not be too early for some FAs. D.O.Riska (FI) agreed with a 'once only' approach and asked for which years are the missing funds needed. Della Negra showed the funding shortfall for construction to be 10 M in 2002, 17 M in 2003 and 16 M in 2004, with M&O costs coming in at 3 M in 2001 rising to 18 M in 2005. The problem starts in 2002 and so must be faced today.

J. O'Fallon (US – DOE) said that discussions with FAs should begin as soon as possible. Although the US will not provide any more money at present, the use of remaining funds and contingency monies can be examined. M. Goldberg (USA – NSF) remarked that FAs will want to know about descoping/staging before considering other options. Cashmore said that would be part of the dialogue. R. Eichler (CH) said that Switzerland would not offer more money and remarked that extra funds are also needed for computing and M&O.

F. Bello (PT) asked whether 44.3 MCHF is the final figure and how urgently does it have to be spent. Della Negra replied that the figure will be revised for October. CMS has a spending profile but the collaboration must know the likely magnitude of the problem soon to know whether descoping is necessary. Because the ECAL crystals are needed now so as not compromise the Higgs search, the 10 MCHF \$/CHF exchange rate increase may have to be accepted, even though it could get better (or worse).

Feltesse said that one solution might be to extend funding over a longer period and asked what the delay (and physics impact) would be if CMS did not do extra work on the surface. A. Hervé (CMS) estimated the delay at 6 months, pointing out that there would be consequences for the other experiments. As for the physics, Ganguli equated delay with disaster given the competition from FNAL. Riska expressed the need for worst case scenarios to highlight the technical/physics consequences. Cashmore added that when the problem falls on one sub-detector resources might have to be moved around as some sub-detectors have higher priority than others. When multiple FAs are involved this gets difficult and it becomes a task for the RRB.

L. Maiani (CERN DG) recalled that the original LHC experiments had been approved with staging/descoping as the only contingency. CMS have made an effort to find out what is missing and <10% at this stage can be regarded as quite good if it is seen as contingency, although the cost of deliverables has also increased. By the end of this difficult year, a clearer picture should emerge of the status of the machine and of computing. Staging and descoping are legitimate but undesirable options with widespread consequences. The DG expressed his hope that the FAs understand that the Laboratory and the experiments are trying hard to assess the present status and the consequences of having 'no money'.

Cashmore reminded the RRB that almost all prices are in 1995 CHF and CMS should address the effects of indexation/inflation on the Common Fund/Projects. The MoU contains a prescription for doing this based on CERN's material Cost Variation Indices.

There being no further comments, Cashmore asked that CMS be encouraged to go ahead with preparing a plan to cover the missing funds needed for completion and to present a progress report to the October RRB.

4. Status of magnet and infrastructure

(Technical Co-ordinator A. Hervé)

(CMS RRB-Tr 2001-105)

Civil engineering

The CMS magnet will be assembled on the surface and then lowered into UX5 to decouple the assembly work from the underground construction work. However, installing and cabling the detector is strongly coupled to termination of the civil engineering. The pillar wall between UX5 and US5 must be finished before cavern excavation can begin. The delay created by problems with the low quality molasse is now expected to be 9 months. CMS has taken that into consideration in a new general schedule v31.

Yoke

Production of the magnet is going well. The barrel yoke and vacuum tank have been delivered, except for the inner vacuum tank. This has been repaired and the two inner rails have been machined; delivery is now expected in May 2001. The first three barrel rings are completed at Point 5. The eight feet and two transport beams have been delivered, completing the Pakistani in-kind contribution, and all feet have been equipped with grease pads delivered by LIP.

The first endcap was delivered from Japan in early April; the second will arrive in September. All endcap carts have arrived, completing the in-kind Chinese contribution. Assembly of the first endcap is starting 23 April 2001.

Coil

For the coil, large procurements have been organized through Institutes, summarised below. Tooling for the swivelling operation is going well in Korea (Hanjung, Pusan) and delivery is expected in September 2001.

<i>Item</i>	<i>Institute</i>	<i>Status</i>
Super-conducting strands	FNAL	OKSC: completed 680 km (for 8 cables); all satisfy specification. IGC: 2 'first article' cables delivered.
Pure Al stabiliser	FNAL	Sumitomo: quality better than specification; delivery complete.
Al alloy AA 6083 reinforcement sheath	FNAL	AluSuisse: now going well; 20 out of 40 lengths produced.
Cabling	ETHZ	Brugg Kabelwerk: 5 good cables produced; 2 lost due to machine fault will have to be replaced.
Coextrusion	ETHZ	Alcatel Cable Suisse: 2 dummy inserts (2x2'650 m) will be used to qualify continuous EB welding line starting end April; 4 real insert lengths produced, 1 for winding prototype, 3 (out of 20) for coil.
Electron Beam welding of reinforcement sheath	ETHZ	Techmeta: commissioning tests of line completed using full section dummy conductors. <i>Milestones:</i> April 2001 (contractual), 1km real conductor; 21 lengths May 2001 to May2002
Winding	INFN	Ansaldo: design and construction of winding line sub-contracted to TPA, Brianza, IT. Machine commissioning starts June 2001 using dummy conductor from Techmeta.

The outer cryogenics system (Air Liquide, France) will be tested in SX5 in 2002 using a dummy load for which metallic structures are presently being procured. An Engineering Design Review EDR for the power circuit was passed in March and procurement can proceed for the power supply, dump resistors, power breakers and current leads. EDRs for the inner cryogenics and the magnet control & safety systems will be held in 2001. This will complete the design and authorize the last major procurements (apart from renting the heavy lifting gantry).

Infrastructure

The general infrastructure is at the final design stage. Integration in the difficult region on top of the coil and around the yoke has been completed. Procurement of the metallic structure that houses racks and services for subdetectors is underway. A special effort is going into active cooling for the sub-detectors and the definition of the gas system. Cooling pipes and gas pipes for the subdetectors will be installed on the yoke starting mid 2002.

Schedules & Milestones

The magnet schedule, 7 months late with respect to TDR, is now based on contractual dates. The delay is still in the shadow of the expected civil engineering delay. Coil project milestones are:

Production of dummy conductor section	September 2000
Order swivelling platform	October 2000
EBW line trial commissioning with dummy section	March 2001
Production of reinforcement for EBW line commissioning and 2 x 2'650 dummy EBW lengths	September 2000
Final engineering of the winding machine	August 2000
First length of real EBW conductor	May 2001
Commissioning of winding machine	June 2001
Delivery of mandrel for prototype	June 2001
Prototype module production	July 2001
Prototype module impregnation	October 2001
Beginning of module production	October 2001
Start cryogenic tests on dummy load	March 2002
End of module production (at Ansaldo)	April 2003
End of equipping cold mass (thermal shields, insulation, swivelling & insertion)	end 2003
Ready for magnet test in SX5	15 March 2004
End magnet test in SX5	30 July 2004

The v31 Master Schedule, for a complete CMS detector ready for the first physics run in August 2006, takes into account realistic civil engineering and magnet schedules. There is a little more time for constructing the magnet but less for cabling and commissioning the sub-detectors. The implications of this will be discussed in the next presentation.

Cost estimate & budget

Including part of the reserve put aside for the conductor, 86% of the budget is committed, 104 MCHF. The cost estimate is maintained at 121.9 MCHF (originally in 1995 prices). It should be known by the next RRB whether there is sufficient reserve left to replace the 2 cables lost due to the Kabelwerk machine fault.

5. Status of integration & detector construction

(Deputy Technical Co-ordinator A. Ball)

(CMS RRB-Tr 2001-106)

Ball outlined the new CMS planning v31. Underground civil engineering work is currently expected to end by 1 April 2004 and 22 months later CMS must be ready for the LHC pilot run. This time interval is regarded as the minimum necessary to have a 'working detector' in place (missing one HCAL endcap, the pixel tracker and part of cabling and trigger/DAQ). There will be about 2.5 usable months during the shutdown after the pilot run, bringing the total assembly and commissioning time to 24.5 months compared with an estimated required 27 months.

Several steps are being taken to compress the underground phase by 2.5 months to have a working detector ready for 1 February 2006. The detector assembly sequence is re-arranged with several activities moved from UX to SX. Negotiations are underway to start pre-installing CMS infrastructure in the UX and US caverns during civil engineering. It is planned to cable and commission first half of detector while the second half is being installed. Ball noted that the time needed for cabling and commissioning has probably been underestimated and may well now be the critical item in completing the detector.

The revised schedule exploits the fact that the surface phase is 4 months longer. More time has been allocated to projects with little time contingency, namely the magnet, forward HCAL and muon system. The magnet test schedule has extended by only 2 months, so the subsequent 2 months have been assigned to activities moved from underground, in particular installation of half of EB and HB which will be lowered as a single unit. As many detectors as possible will be pre-cabled and pre-tested on the surface, requiring services and local DAQ in the SX building not foreseen in the original planning. Installation tooling has to be designed with surface or underground use in mind and precabling as much as possible on surface requires more patch panels/disconnects.

In case of further delay in completing the civil engineering, it has been decided to keep sub-detector delivery target dates more or less unchanged and to move even more activities to SX. It will be difficult to compress the time needed for cabling, commissioning and installing the trigger/DAQ. Few of these activities can be moved to the surface and they may define the minimum underground phase attainable. Work is going on to understand the procedures in more detail.

For each sub-detector the cumulative progress of critical items are being monitored and measured against target curves matched to v31 schedule:

Tracker	module assembly (not yet in bulk production)
ECAL	crystal delivery (using large ingots) and supermodule/dee assembly
HCAL	optical megatile manufacture
Muon Barrel	drift tube chamber assembly
Muon Endcap	cathode strip chamber assembly

At present subdetector manufacturing is on track to meet v31 targets. Based on schedule v26 of August '99, 89% of the milestones to be done by 1 April '01 were completed. The milestones will be re-baselined to v31 by June, and 50 or so new ones will be added.

6. Report from the LHCC (LHCC Scientific Secretary E. Tsesmelis)
(CMS RRB-D 2001-103; CMS RRB-Tr 2001-103)

The presentation covered LHCC sessions in November 2000, January 2001 and March 2001.

Although some critical items have been identified, the LHCC commends the impressive

progress the CMS Collaboration has made towards being ready to record data at the beginning of LHC operation in 2006. The proposed detector lay-out at LHC start-up will allow CMS to satisfactorily address the physics issues for Higgs and SUSY.

The LHCC noted that civil engineering delays have serious implications for detector installation and commissioning and may lead to more resources being needed to carry out additional tasks on the surface. It also noted that detector installation is foreseen beyond the initial running period.

Reviewing the CMS sub-systems, good progress was noted in the Tracker performance studies and its alignment system, in the ECAL silicon APDs and crystal production process and in the development of the HCAL absorber and scintillator. However, difficulties in the QIE and HPDs of the HCAL were found to be either on or approaching the critical path. Likewise, the HF detector is now on the critical path. The Committee considers that CMS is making every effort to complete the HCAL and HF in time for agreed installation milestones.

Critical items for the Muon Spectrometer were identified as the barrel DTs and the electronics for the CSCs. The Committee considers that the extra capacity being set up in Torino will help recover the 6-month delay in the barrel DTs and that CMS is trying hard to complete the CSCs in time to meet installation milestones. The LHCC is currently evaluating the performance and ageing properties of the RPCs of all LHC experiments and will present its conclusions and recommendations at its May 2001 session.

Concerning the CMS Level - 1 Trigger, the Research Board concurred with LHCC recommendations for approval of the TDR, pending a final cost evaluation.

The LHCC recently performed a cost review of CMS and noted that CMS is following strict budget control, and that the cost overrun of $\approx 6\%$ (excluding funding shortfall) is very reasonable for a project of this scale. An LHCC Comprehensive Review of CMS will take place in October 2001.

The LHCC is currently reviewing the data management and computing requirements of the experiments. The Committee puts fulfilment of these requirements on an equal footing with the successful construction of the LHC machine and the detectors for the success of the LHC physics programme. It will present its conclusions and recommendations in May 2001.

Discussion

Asked by Ganguli about the ATLAS cost overrun, Tsismelis confirmed that it is also $<10\%$. The LHCC does not find these levels unreasonable and notes the collaborations' considerable efforts to contain costs.

7. Financial matters (*CERN Finance Division Leader, A.J. Naudi*)

Status of collaboration accounts

(CMS RRB-D 2001-100)

Updating the distributed document, Naudi announced further payments, in kCHF:

IN2P3	760
Croatia	5
Turkey	35
Helsinki (paid in full)	300
Estonia	10

plus a substantial amount from US-NSF. Expenditures increased by 264 kCHF to reach ~ 500kCHF and outstanding commitments to ~10MCHF. Invoices still outstanding for the following countries: AT, BE, CH, DE, GR, HU, IT, ES and PL amount to some 1.5 MCHF. Naudi urged FAs to pay as soon as possible to face the calls on the Common Fund and to accumulate interest (interest had amounted to 610kCHF by the end of 2000).

Naudi announced that at the next RRB he will present the findings of CERN's External Auditors, the Tribunal de Cuentas (ES), appointed by Council, who will submit their report towards the end of May.

- *Discussion*

H. Borns (AT) asked that delegates be given paper copies of the updates at the meeting. Naudi said he would do his best to satisfy this request.

Summary of market surveys & tenders

(CMS RRB-D 2001-101)

Naudi reported the following updates to the distributed document:

IT-2731/EP contract sent to Erbsloh Aluminium (DE) 21.03.2001

IT-2811/EP letter of intent sent to Ericsson Cables AB (SE) early in April 2001

MS-2690/EP and MS-2878/EP: the issue of the Call for tenders will be May 2001 instead of March 2001.

8. Budget matters

(Resources Manager D. Blechschmidt)

Summary of Expenditure 1995 - 2000

(CMS RRB-D 2001-98, CMS RRB-Tr 2001-98)

Blechschmidt announced that 55.5 MCHF were committed and 58.7 MCHF paid during 2000. Of the total cost estimate of 456 MCHF (figures in the original MoU and the Si tracker amendment to the MoU), 39% has now been committed and 28% paid during 1995 – 2000:

Overview of commitments/payments for 1995 – 2000:

	<i>Cost Estimate</i>	<i>Committed</i>	<i>Paid</i>
Magnet	121.8	86%	53%
Tracker (full Si)	77.6	3%	2%
ECAL	85.7	26%	15%
HCAL	41.8	52%	52%
MuonDetector	60.8	37%	35%
Trigger/DAQ	37.5	5%	5%
Offline Computing	3.6	23%	22%
Infrastructure	27.2	15%	11%
<i>Totals</i>	<i>456 MCHF</i>	<i>39%</i>	<i>28%</i>

	<i>Committed</i>	<i>Paid</i>
Magnet CP	104.6	64.1
Offline CP	0.8	0.8
Subdetectors	74.5	62.5
<i>Totals</i>	<i>179.9 MCHF</i>	<i>127.4 MCHF</i>

Commitments/payments for Common Projects:

	<i>Committed</i>	<i>Paid</i>
Common Funds	34.4	24.2
Contracts (placed by Institutes)	69.2	39.3
In-kind (China and Pakistan)	1.8	1.4
<i>Totals</i>	<i>105.4 MCHF</i>	<i>64.9 MCHF</i>

Blechsmidt showed the levels of Common Project commitments (as due under the signed MoU) currently reached by each FA and encouraged FAs with low levels to pay as soon as possible. (Similar levels were presented later for sub-detector commitments and for CP and sub-detector payments).

Discussion

J. Sacton (BE) remarked that FAs only pay bills that they receive. Blechsmidt noted that the rate of invoicing needs to be accelerated. Cashmore added that, whenever it is convenient, FAs should let the collaboration know if they are ready to pay earlier.

Commitments/payments for Common Funds:

	<i>Committed</i>	<i>Paid</i>
Income from FAs	30.7	30.7
Accrued interest	0.61	0.61
Total Income	31.3	31.3
Total committed	34.4 MCHF	24.2 MCHF
<i>Carried forward to 2001</i>	<i>-3.1 MCHF</i>	<i>+7.1 MCHF</i>

Commitments/payments for sub-detectors:

	<i>Committed</i>	<i>Paid</i>
Tracker	2.5	1.5

ECAL	22.2	13.21
HCAL	21.6	21.64
Muon detector	22.3	21.23
Trigger/DAQ	1.8	1.81
Infrastructure	4	3.11
<i>Totals</i>	<i>105.4 MCHF</i>	<i>62.5 MCHF</i>

Trends of accumulated commitments show them starting to slightly exceed original estimates; payments, which were lagging, are now converging with the planned profile.

Preliminary Draft Budget for 2002 (CMS RRB-D 2001-99, CMS RRB-Tr 2001-99)

In 2002, the magnet and all sub-detectors will be in construction. Preliminary estimates for payments add up to ~ 90 MCHF, ~ 20% of the total estimated construction costs. These figures are tentative; they have not been discussed or agreed by the FAs.

Magnet	17	10.2 from Common Fund
		6.9 payments to contracts
		0.1 in-kind
Tracker	22	1 pixels
		21 Si detector
ECAL	22	6 crystals
		9 electronics
		5 structures & assembly
		2 preshower
HCAL	8	3 barrel HB
		1 outer barrel HO
		2 endcaps HE
		2 forward HF
Muon detector	13	4 Drift Tubes
		5 ME1/1 & endcap CSCs
		3 barrel & forward RPCs
		1 alignment system
TriDAS	3.5	3.3 Trigger
		0.2 DAQ
		0.1 detector controls
Offline Comp.	0.5	
Infrastructure	4	1 access & survey
		1 installation
		1 cooling & ventilation
		1 safety & shielding
<i>Total</i>	<i>~ 90 MCHF</i>	

If approved, total payments will reach 292 MCHF by end 2002, ~ 64% of total. Blechschmidt noted that under the present assumptions (carry-forward from 2001 and income and expenditure in 2002) the CMS Common Fund will have a shortfall of ~ 5 MCHF in 2002. FAs are invited to consider advancing payments or underwriting formal guarantees for future payments of remaining Common Fund contributions.

I.F. Corbett (UK) correctly presumed that these preliminary figures make no assumptions on cut-backs or descoping to handle the currently projected total shortfall of 44 MCHF and asked that the October budget be consistent with any actions taken. Cashmore

remarked that CMS are already doing a certain amount staging and that the bigger the CF the easier it is to handle the cash flow. He also noted that FAs paying into the CF early will not have to worry about inflation indices.

9. Discussion on LHC Computing Review Report (RRB-D 2001-03)

At Calvetti's suggestion, Cashmore started by summarising the discussion that had taken place at the LHCb RRB earlier in the day. It is recognized that the computing challenge is such that it cannot be met by a fragmented approach. CERN should take the lead in preparing an integrated plan for discussion at the October RRBs. The FAs obviously want to be involved and want to know what resources are needed both in CERN-IT and in the experiments. The CERN Tier0+Tier1 facility must be a common project between the experiments, with CERN-IT playing a strong role backed by adequate resources. Discussions took place in March with the SPC and Committee of Council on the basis of a paper 'Building the LHC Computing Environment', and CERN is now developing a plan for discussion during the June Council week. A two-phase project is foreseen, 3 years for setting up the GRID and prototype infrastructure and 3 years for setting up the final LHC computing environment. Costing the second phase is necessarily based on extrapolation. Some decisions are urgent but there are still many industrial cycles left before LHC computing proper will be underway.

Concerning manpower levels, H. F. Hoffmann (CERN Director for Scientific Computing) informed the RRB that, as well as CERN-EP scientific programmers, 40 – 50 CERN-IT programmers are working on LHC computing and 10 additional posts have been allocated recently. There are also 10 EU–DataGRID people at CERN. In the US the NSF and DOE have proposals for GriPhyn and PPDG projects similar to DataGRID. The plan currently being developed aims to define the supervisory structure for a coherent project that must include external Tier1 (and possibly Tier 2) centres. The first of the two phases mentioned by Cashmore starts now and involves mainly software and middleware, needing substantial manpower but relatively modest material investments. It should consist of tests of increasing size, with a final prototype of $\approx 50\%$ the complexity of one LHC experiment. This phase will have to be covered by IMoUs or software agreements.

Ganguli asked about sharing the total hardware cost of ≈ 240 MCHF, and whether networking costs are included. Hoffmann explained that cost ratios are $\text{Tier0}:\Sigma\text{Tier1}:\Sigma\text{Tier2} \approx 1:1:1$. He added that the 1/3 to be paid by CERN is not in current long-term budget plans. Networking costs are not included. It is assumed that they continue at present levels and that capacity will double each year at constant cost. Riska asked whether network throughput doubles each year. M. Delfino (CERN-IT Division Leader) said that because of a problem, present loads only reach about 10% of available bandwidths. Eichler asked what would be the effect on the physics if CMS produced only half the amount of data thereby reducing the computing requirements. Cashmore commented that the LHCC may eventually have to discuss such options but at a second phase after more prototyping. Della Negra pointed out that the required computing capacity is not directly linked to trigger rates and physics; the facilities are also used for activities such as simulation. He also noted that the 240 MCHF estimate has a large error

bar and only covers phase 2. It is vital that LHC data can be distributed to allow all collaborators to profit from the experiments and so it is urgent to go ahead with phase 1 prototyping.

Corbett said that CMS is facing a ≈ 40 MCHF funding problem, ATLAS likewise, CERN cannot provide its part of the LHC computing, M&O funds are needed, and so on. PPARC does not have the money to meet these costs, presumably also the case for other FAs. The RRBs are the only place to face all of these problems generically across the experiments. These items will dominate the RRBs agendas for some considerable time.

There being no further comments, the meeting moved to questions of CMS-specific computing in the next presentation.

10. CMS IMoU for Core Computing & Software
(*Deputy CCS Project Manager D. Stickland*)

Several documents associated with this presentation can be found at:

<http://cmsdoc.cern.ch/cms/cpt/april01-rrb>

The LHC distributed computing environment is hard to simulate. A series of tests of increasing complexity is needed to reveal and resolve successive levels of difficulties. To achieve this, two major problems must be addressed urgently, the acute lack of software engineers and the need for prototyping. CMS also needs massive production runs during 2001-2004 for trigger, detector and physics optimisation. Stickland noted that LHC simulation and reconstruction is ≈ 200 -400 times more complex than for CDF/D0. To meet these challenges, CMS is preparing a Computing IMoU for the end of 2001. The participation of FAs is strongly encouraged and a mechanism will be established to account for contributions. The full LHC Computing System for CMS will be covered by an MoU that should come into effect ~ 2004 .

The CMS software project, CPT, has three closely coupled components, Computing and Core Software, Physics Reconstruction and Selection, and TriDAS (online farm and software). A plan is in place that allows clean and co-operative management of physics software developed by the Institutes, CMS core software developed by the collaboration, and central core software provided by CERN/IT.

According to March 2001 estimates, CMS manpower needs for on-line and off-line core software, in FTE software professionals, are:

2000	2001	2002	2003	2004	2005	2006	2007
26	28.5	35	39	41	42	43	40

Despite recent increases in effort from Italy, France and Russia, about half of the required manpower is currently missing. No institutes are formally responsible for core software,

so it is taken to be a collaboration-wide responsibility. Based on a recent count of CMS scientists, 42 FTEs in the March 2001 planning for 2007 would be provided *pro rata*:

	<i>Scientists</i>			<i>FTEs</i>			
USA (DOE & NSF)	262	23.1%	9.7	Korea	17	1.5%	0.6
Italy	190	16.7%	7.1	Bulgaria	17	1.5%	0.6
CERN	107	9.4%	4.0	Greece	15	1.3%	0.6
RDMS-DMS	72	6.3%	2.7	Austria	15	1.3%	0.6
RDMS-Russia	60	5.3%	2.2	Finland	12	1.1%	0.4
France	59	5.2%	2.2	Portugal	10	0.9%	0.4
Germany	45	4.0%	1.7	Switzerland (PSI)	8	0.7%	0.3
UK	38	3.3%	1.4	Poland	8	0.7%	0.3
Switzerland (ETHZ)	36	3.2%	1.3	Turkey	7	0.6%	0.3
Spain	34	3.0%	1.3	Pakistan	7	0.6%	0.3
China (NSFC)	33	2.9%	1.2	Taipei (NSC)	6	0.5%	0.2
India	25	2.2%	0.9	Croatia	4	0.4%	0.1
Belgium	25	2.2%	0.9	Estonia	2	0.2%	0.1
Hungary	20	1.8%	0.7	Cyprus	2	0.2%	0.1

Fractional FTEs can either be increased to one person or replaced by a financial contribution.

Before purchasing the production system, the aim is to double the complexity (number of boxes) of common LHC prototypes each year to reach 50% of one experiment in 2004. CMS will try to match the computing challenges to physics and detector milestones. Stickland noted that there will be a DAQ TDR end 2002, a computing TDR in 2003 and a physics TDR in 2005.

To test the distributed computing model and to build up relevant experience in regional centres, CMS needs 30% of the CERN Tier0 (either time-shared or by capacity division), 2 effective regional Tier1s and 10 effective Tier2s. The plan is to have a 20% data challenge in 2004 of 40M events, reconstructed, distributed and analysed at a rate of 20Hz for 1 month.

The following costs in MCHF are anticipated for CMS LHC computing:

	<i>Not included in Computing Review</i>			<i>Included in Computing Review</i>			
	<i>2002</i>	<i>2003</i>	<i>2004</i>	<i>2005</i>	<i>2006</i>	<i>2007</i>	
<i>Hardware costs in MCHF</i>	<i>Prototyping</i>			<i>LHC Computing</i>			<i>Totals</i>
30% T0/T1 at CERN	0.8	0.8	1.5				3.1
T0/T1 LHC computing for CMS				5	5	6.7	16.7
2 effective regional T1s	1.3	1.3	2.6				5.2
5 effective regional T1s				7.4	7.4	9.8	24.5
10 effective regional T2s	1.4	1.4	2.9				5.7
25 effective regional T2s				8.5	8.5	11.3	28.3
<i>Totals</i>	<i>3.5</i>	<i>3.5</i>	<i>7.0</i>	<i>20.8</i>	<i>20.8</i>	<i>27.8</i>	
	<i>IMoU 14.0 MCHF</i>			<i>MoU 69.4 MCHF</i>			<i>83.4</i>
<i>Core S/W Personnel FTEs</i>	<i>Prototyping</i>			<i>LHC Computing</i>			
Offline	25	34	35	35	36	36	
Online	4	6	8	9	10	6	
<i>Totals</i>	<i>29</i>	<i>40</i>	<i>43</i>	<i>44</i>	<i>46</i>	<i>42</i>	

It is supposed that prototyping facilities can be optimised so that the cost for 4 experiments is less than 4x the cost for CMS. Stickland also noted that CMS already has more Tier2s operational than expected because GRID R&D money is available and is being used. It is hoped that this kind of funding will continue to provide resources, particularly during prototyping.

The size of the facilities is designed to match the computing required at expected LHC luminosities. Taking CMS data at 100Hz, the Tier0/Tier1 CERN duo will have to reconstruct 200 M events per month with 100 M events re-reconstructed at off-site Tier1s. With these loads, the computing facilities foreseen in 2006 are sufficient but not luxurious. Among Tier1 candidates are centres in US/FNAL, Italy/INFN, France, UK, Germany and a Russian-cluster Tier2 cluster of Tier1 capacity. These are not all the same size and some are shared with other LHC experiments.

Closing his presentation, Stickland reminded the RRB of some of the Computing Review recommendations, notably the construction of a common prototype as a joint project of the four experiments and CERN/IT, along with the major Regional Centers. CMS believes a multilateral IMoU is the best approach to this challenge, starting the process towards a full MoU for LHC Computing.

Discussion

Stickland confirmed Ganguli's observation that engineers are not in the head-count used to estimate participation, and added that at this stage the model was for guidance only. Corbett noted that an expected 20 MCHF per year will be spent on hardware. Personnel costs are typically 3 - 4 times hardware costs, whereas the core FTEs mentioned account for only some 10 MCHF. So where are the other 100 or so FTEs? Stickland replied that Tier1s use a lot of manpower and many are being set up in existing computing centres.

Commenting on the need for prototyping, Delfino noted that the task of managing 20,000 interworking devices is unprecedented and cannot be tackled using traditional methods. Stickland added that current prototype LHC computing is two orders of magnitude away from final estimates in almost all parameters.

G. Barreira (PT) asked about small countries that will have only Tier2s and Tier3s and will have to connect to an expensive Tier1 centre. Will they have to contribute to Tier1 capital investment or running costs? Cashmore said that the GRID-user community is still trying to understand this kind of question. Hoffmann remarked that acceptable ways have been found for sharing financial responsibility for complex, common property such as detectors and that phase 1 of the LHC computing project will be used to work out fair ways to balance computing costs. Della Negra remarked that such issues will have to be addressed in computing MoUs. The common LHC project will have a global, common part with own management structure and funding, and separate agreements for each experiment. So, for example, the arrangement for Portugal will have to be spelled out in the CMS MoU. Barreira pointed out that Portugal is in ATLAS as well as CMS so co-

ordination is needed. Cashmore added that there is an overall need for a coordinated plan because CERN's Tier0/Tier1 facilities and many Tier1s will be shared.

Asked by Barriera what would be in an IMoU, Cashmore stated that is would define mainly the CMS software engineers needed to complete the plan for prototyping with Tier0/Tier1 and central GRID facilities at CERN. Stickland remarked that the FTEs mentioned are not building the GRID, they are integrating CMS software with the GRID.

Ganguli asked about networking costs and Cashmore said that it has been assumed that countries will keep up anyhow. Hoffmann said the Review had considered the situation in deregulated countries and concluded that other countries will probably follow by ~2005, bringing down prices. For example, networking costs in the EU are now approaching US levels and have previously been some 10 times higher.

S. Bhave (IN) said that if an IMoU is to be ready for signing by end 2001, items need careful costing. Cashmore agreed that CMS will have to make best estimates and that end 2001 may be an optimistic target; nonetheless an important first step is underway. If India could supply a software engineer, here or at home, or perhaps 2 working together at home, that could satisfy India's obligations.

Cashmore mentioned that CERN is trying to access new funds via the general interest in the GRID, emphasising the idea that the GRID will happen if the focus is on LHC computing. He concluded the discussion by underlining the need for the experiments and CERN to produce a plan for providing, in a balanced way, an LHC computing environment across all collaborating countries, institutes and FAs.

11. Discussion on M&O paper (RRB-D 2001-04)

As in the computing discussion, Cashmore started by summarising earlier discussions in the LHCb RRB. It is very important to identify the total and real costs of running the LHC experiments. Present figures in the M&O paper are indicative rather than accurate. Items need tighter definition and categories need further consideration. Some costs may already be covered in the construction MoUs. Also, cost estimates have to be revised in view of the new LHC start-up date of 2006.

Cashmore will set up a small group, consisting of the 4 resource coordinators and 3 or 4 members of CERN management, to produce clear definitions of the line items, definitions that will be included in the M&O MoU. With the help of the RRBs, he will also set up so-called scrutiny groups to examine the collaborations' cost figures and guarantee their integrity. The scrutiny groups will have 3 or 4 members nominated by the RRBs for each experiment, and 3 or 4 nominated by CERN common to all experiments to provide uniformity. Cashmore asked delegates to send names to him by **11th May**. This interim procedure is necessary as the October RRBs will have to address M&O funding for 2002. The formal procedure for setting up scrutiny groups will be defined in the M&O MoU.

A draft M&O MoU will be prepared by CERN and the experiments for discussion in October. Cashmore asked delegates to send their input to him by **25th May**. Clearly RRB discussions are needed before the October meetings and 2 or 3 dates in September will be proposed, hopefully allowing all delegates to attend on at least one day.

FAs have made it clear that they need to know the level of rebates when considering A-type costs. This has to be taken up by the CERN Directorate. Concerning B-type costs, there is a consensus that these should remain with the collaboration and the FAs.

Borns remarked that many costs previously C-type now appear as A-type and hoped that the 'items working group' will clarify several outstanding points. She added that people in Ministries are faced with these tables for just a few hours per year and need convincing that all costs are necessary and that there is no double counting. Cashmore agreed on the need for clarity and commented that rebates, a mechanism already used at DESY, will solve part of problem of A-costs. In response to some particular points raised by Borns, he said that the 'infrastructure' discussed by the CMS resource manager is at the pit, just as the computing in the M&O paper is at the pit and is thus distinct from the computing handled by the LHC Computing Review. He remarked that 'new costs' are becoming evident, costs to completion, costs for M&O and LHC computing. It is the RRBs' job to sort out what is possible and set priorities.

Calveti asked where the numbers come from and whether delegates could be given more details before before 25th May. The numbers come from the collaborations and Cashmore offered to try to provide more details.

Wormser applauded the SPC recommendations presented during the plenary session but wondered how many have been followed. Cashmore replied that 'best' practices at LEP and HERA have been followed. For example, attributing A-costs by qualified scientists, now widely preferred, was not done at LEP. Wormser noted that the question of A/C categorisation is closely linked to rebates and that FAs need to know more about these. He also found it hard to understand the time evolution of total costs, flat between 2002 and 2004 followed by a steep rise. Cashmore stated that after clearer classification, the experiments must re-estimate their costs for examination by the scrutiny groups.

Replying to Ganguli's query as to why the idea of A-cost sharing by a 50/50 split scientists/costbook had been dropped, Cashmore said that sharing by qualified scientists (excluding engineers and students) was favoured by both the SPC and a substantial fraction of the community. He reminded the meeting that the collaboration decides on the sharing algorithm for B costs.

O'Fallon noted that M&O costs have been faced before, for example at LEP; after a struggle to understand them, the requests are 'scrubbed' and finally funded. However, serious amounts of money are now involved and it is important that the process is transparent; the basis for the costs and for their distribution must be clear. US agencies have the difficult task of convincing Congress that the requests are legitimate and not simply a response to a bill that arrives because they have signed up to an experiment.

G. Vesztergombi (HU) felt that sharing A-costs by scientists gives countries the possibility to cheat by withholding names from publications. He said that the algorithm could cause serious problem for some small countries. For example, Hungary, a Member State, pays 0.5% of the CERN budget; in CMS it pays 0.2% of the costbook total and would pay 2% of the running costs if shared by scientists. Cashmore repeated the consensus view that scientists benefit from the output of the LHC and so should bear the financial responsibility. Vesztergombi noted that skilled scientists contribute to the collaboration and that a 50/50 split would help poorer countries. Bhave suggested that each country will favour the scheme that benefits it most and a 50/50 split could be a compromise. Cashmore said that part of the A-costs could be paid in-kind. Calvetti added that this kind of scientist/costbook disparity is taken into account by the fact that A-costs are only a fraction of the total; the B-costs are probably comparable.

Feltesse asked whether manpower is included in the figures. Cashmore said that there is some and that it must be clearly identified and properly costed so that payments can be made in-kind, which could help some countries. Calvetti commented that the RRBs need to know the order of magnitude of problem for particular countries. Riska said that given the strong reservations about sharing by scientists, the RRBs need tables showing the consequences of applying different algorithms. D. Green (USA-DOE) enjoined the meeting not to forget B-type costs.

Maiani then expressed his views. The substantial costs for operating LHC detectors are consequent on decisions concerning those detectors and are not directly determined by CERN. Therefore operating costs must be discussed by the same body that discusses detector construction and with a uniform approach to all experiments. The scientists in the collaborations bring scientific returns and benefits to their countries and so the number of them would appear to be a natural parameter for sharing general costs. If this approach creates distortions, remedies can be sought for pathological cases. Maiani found it difficult to believe that countries would remove their scientists just to pay less. He appreciated that Ministries will question M&O costs, some feeling that they pay them already. However, at the LHC, the set of participating countries is much larger than the set of supporting countries. He deemed it very useful to distinguish running the laboratory from running the detectors, adding that countries participating in experiments should participate in expenses. The ICFA guidelines for an open laboratory must be followed, but CERN Member States are no longer in a majority hence the concept of rebates as a way to achieve equitable cost sharing; the contributions of some non-Member States to CERN will also be recognised. Maiani concluded by underlining the importance of knowing the costs of running the detectors, costs that are not within CERN's control.

11. Summary and future activities (R. Cashmore)

Good progress has been reported on common items and sub-detector systems and CMS will undergo an LHCC Comprehensive Review before next RRB.

CMS is very grateful to the RRB's for giving the go-ahead for Bakelite purchase by

borrowing from the Common Fund.

The estimated cost to completion of CMS will now be followed up in detailed discussions with the FAs with a progress report given to the October RRB. Half way through the project a 6 - 10% overrun can be viewed as quite an achievement. Of course, the collaboration must try to reduce this before asking for more resources. FAs are encouraged to pay promptly into Common Fund to to minimise cash flow problems.

There is work to be done to get a coherent approach to LHC Computing. On M&O, vital for the experiments, lots of homework is needed before October.

Proposals for scrutiny group members should be sent to Cashmore by May 11th and comments concerning preparation of the draft M&O MoU by May 25th. Delegates will be informed of dates of September interim meetings as soon as possible.

12 Dates of next formal meeting

October 22nd/23rd 2001

The format of the next meeting may be recast to include plenary discussions of generic items. Thereafter, for 2002, it must be decided whether a 2-day format is sufficient to allow the RRBs to conduct their business satisfactorily.

Cashmore informed delegates of an LHC Symposium, aimed at young physicists, to be held in Sardinia, **25th – 27th October 2001**.

13. Any Other Business

There being no other business, the Chairman closed the meeting