CMS Response to PPRP Visiting Panel feedback

1 The Panel would like more information on certain descope scenarios:

(a) Using the model of 30% cuts outlined in the previous response, which passes significant risk to STFC, what can be done to reduce the risk, and what are the consequences?

(b) Using the 30% descope as the starting point, what funding should be added to maintain the deliverables and reduce the risk to a reasonable level?

(c) What would be the prioritisation if a 50% cut was required, including the political and experimental impacts?

1 (a) Using the model of 30% cuts outlined in the previous response, which passes significant risk to STFC, what can be done to reduce the risk, and what are the consequences?

Although the time available has been short, we have looked at the resources needed and the capacity to absorb a 30% reduction in new funding, while reducing the risk. We have looked more carefully at the staff effort required and can see somewhat better ways of optimising the staff usage, rather than the blunt approach of removing one or more RA posts in all institutions. We have also looked at the working allowances (WAs) and tried to see if we can target some items more precisely to help reach the cost reduction required.

For example, the WP3 MP7 board had not been delivered at the time of the proposal, and by the PPRP meeting only the second of two initial boards had been tested, having found assembly faults with the first. More detailed tests have since been possible and the board is working well, with only apparently minor issues to solve. It might therefore be possible to reduce some of the working allowance for WP3. However, the FPGA still burns more power than expected, according to the manufacturer's specifications. This is under discussion with Xilinx and is a concern. In addition, the 72 link version of the Virtex7 FPGA will not be available until about March 2013, so there remains an uncertainty about cost, power and quality, which might affect the schedule and overall system design, e.g. requiring more modules. On these grounds, it does not seem wise to reduce the working allowance too much. However, we have further reduced the firmware sub-contract provision to £135k.

The WP4 pixel FED board is similar to the MP7, but should use a less powerful and less expensive FPGA. Having met the demanding requirements of the trigger processor with the MP7, the FED looks somewhat less challenging. However, it must be constructed in a modular way, with FPGA mezzanine cards (FMCs), and care is needed for the deserialization and error coding, as well as evaluating new optical receivers. The Token Bit Manager chip which will deliver the signals to the FED has been redesigned but not yet delivered and there have been discussions about more than one variant. We believe we must preserve sufficient contingency to accommodate delays and limited changes in requirements.

In WP2 the main non-staff cost is the ASIC manufacturing cost. How much working allowance should be allowed, and can the cost or risk be shared? The proposal allowed for one full engineering run (providing full wafers) and a half run, assuming that another customer could be available to share the cost. In the R&D we have been fortunate that RAL TD have been developing a large ASIC for the DESY XFEL project and our submissions have been able to coincide, which has been extremely helpful to both of us. It is not clear if this can continue, as the XFEL chip is close to final, while the CBC envisages another stage after the CBC2. We also allowed for another half an engineering run as working allowance. In the 30% cut scenario we reduced the cost to 1.25 engineering runs, with 0.25 engineering run as working allowance, although the precise division between the main activity and working allowance is somewhat arbitrary.

There have not been many Multi-Project Wafer (MPW) customers with whom to share manufacturing costs. The alternative of a MOSIS commercial MPW is expensive (~$200k) and provides only very few (~40) cut die, with no possibility of obtaining full wafers. It could be considered for a run aimed at final fine tuning of the design, and would leave the cost of the final production masks to be borne by the CMS Tracker project. However full wafers will be required during the prototype phase to build the wafer test system and for commercial assembly, as well as ensuring sufficient working chips for prototype modules. So the overall budget of 1.5 engineering runs cannot be further reduced. During the CBC design phase, from experience, we do not expect other CMS collaborators to step forward with offers of contributions. It might happen for the final pre-production submission.

We have completely removed the contribution to module construction costs.

The conclusion on the working allowances for the capital items (FEDs, trigger boards and ASICs) is that there is little to be gained by further WA reduction and no way to decrease the risks except by making a larger financial provision.

For the staff effort, we can organise the effort differently, which has some benefits, but only at the expense of removing staff elsewhere. For example, it is desirable to maintain the RA effort in Brunel, which levers out some additional university funded resources as well as CG-funded academic effort. The Brunel RA will now contribute to offline software in both WP4 **and** WP3. To make the best use of RAs and save money, one or two other RAs can be shared between WP3 and WP4, which we hope can profit from the different profiles, as well as ensure that common software is utilised as much as possible.

However, it seems necessary to substantially reduce further the request for TD effort on WP3/4 (not on ASIC design for WP2) to free funds for RAs. This is essential because a significant amount of effort will be required in CERN during commissioning of the trigger and the pixel system, and this must be preserved. We believe that most of the firmware will be developed in the early phase of the projects and longer term maintenance by RAL TD will be limited. In contrast, we will need firmware effort in CERN and online software for firmware validation and commissioning, for which flexibility and fast response are essential, which will be provided by the RAs.

Small savings are possible by assuming realistic start dates for new posts after about 4 months, since the recruitment process usually introduces such delays, and being as tough as possible on institute equipment, most of which is for μTCA crates and infrastructure and wafer and module test equipment.

The conclusion on staff effort is that we can manage the risk in the best way possible by deploying our existing engineering and RA effort but we cannot reduce it significantly within the 30% cut scenario.

Thus, the revised budget aiming at 30% reduction is similar to what was explained in the original response, arriving at a new total cost of £8.21M, compared to £11.49M in the proposal; a reduction of 29%. With the savings described above it is possible to gain a small amount of effort (72.5 FTE instead of 71.2 FTE, before the loss of 3 SY of Fulcher CG effort) but the WA remains at £385k, and there is no explicit staff WA in the budget.

The only alternative to the scenario above would be to cut completely one WP. Why we do not believe this is a viable option is covered in the answer to Q1(c). The arguments are essentially the same for the 30% reduction scenario as for a 50% reduction, though the impact of the larger cuts would certainly be even worse.

In summary, except by cutting a WP completely, with the impact on UK standing, reputation and leadership described in the answer to Q1(c), it is extremely difficult to reduce the risk. The funds allocated to the WA are reduced from about 9% of the proposed original budget of £11.5M new money, to less than 5% of the reduced budget of £8.2M, and with a 30% reduction in new staff.

8 November 2012

1 (b) Using the 30% descope as the starting point, what funding should be added to maintain the deliverables and reduce the risk to a reasonable level?

We have started from scenario outlined in the previous answer and added the most crucial activities.

The new total cost in this model is £9.63M, compared to £11.49M in the proposal; a reduction of £1.85M, or 16%. This was achieved by taking the following measures, all of which aim to reduce the risk to deliverables as their primary motivation

- Restoring the working allowances of WP3 and WP4.
- Restoring ASIC fabrication costs of £534k, but leaving the working allowance at £123k.
- **Not** restoring the contribution to module construction, saving £89k compared to the proposal.
- Restoring the project students to the original proposal level, which ensures continuity in expertise as well as an inexpensive means of providing several FTE.
- Restoring new staff costs to £6.35M (including students) which provides 88.2 staff years of new effort, taking into account the 3 SY lost from CG provision for J Fulcher, compared to 101.2 SY in the original proposal and 69.5SY in the scenario of 1(a).
- Restoring infrastructure costs of £100k to WP3, which can be used either to share the costs of installation of the trigger optical hardware, or for the next generation developments at a reduced level compared to the proposal.
- Increasing the specialist sub-contract effort, for firmware and board design, back to the reduced value, compared to the proposal, of £160k.
- Restoring some of the RAL TD funding for non-ASIC activities, i.e firmware, board procurement and acceptance, to a total of 4.5 SY, including 1 SY for working allowance. This can be profiled to match requirements, so can be used carefully. The original proposal requested 9.6 SY for these staff. However, much of it came late while we require most of it in the early years, and the RA effort is more flexible for deployment at CERN.
- The total working allowance (WA) is now £558k, which represents 6% of the total new budget, but with the flexibility to be cautious in the timing of commitments, e.g. for TD staff use and ASIC fabrication, so as to ensure a sufficient extra margin.

We would still point out that the original proposal was not overgenerous in provision of effort and working allowances and that in large projects there is always a danger that requirements will change quickly. When the proposal was written, the 125 GeV boson had not been discovered and since then the urgency of upgrading the trigger has increased, with much increased emphasis on being ready for worst case conditions in 2015. As the realisation of the importance of low trigger thresholds has dawned, there has also been dramatically heightened interest in increasing the CMS trigger rate and latency to, e.g, 1MHz and 25μs. Both of these have significant implications for the CBC developments where, again, reasonable provision must be made for the unexpected.

1(c) What would be the prioritisation if a 50% cut was required, including the political and experimental impacts?

The answer to this can be formulated in a number of ways and would involve a lengthy discussion with the rest of CMS to agree priorities. It clearly could not be achieved without cutting out one of the WPs entirely, which would then leave CMS dangerously exposed. It is probable we would not get a rapid answer from CMS since each sub-detector project would judge that the removal of our contribution would be extremely undesirable and therefore would argue that the cut should fall elsewhere.

Actually a reduction of funding of this magnitude would be far worse: it requires a cut of £5.75M to be made. This could only be achieved by cutting **more** than a single work package. Removing WP2 and WP3 together would slightly exceed the savings required, while cutting WP4 alone would be insufficient. So the choice would be stark: **to achieve a 50% cut only one WP could be retained unscathed**. Both WP2 and WP3 would be lost if the first option were chosen, while the second option of eliminating WP4 would still require a further £1.1M to be removed from the budget.

Scientifically, cutting both WP2 and WP3 would be devastating to the UK contributions to CMS upgrades, leaving us only with an important role in the pixel DAQ but discarding most of the benefits of the investment in R&D in recent years. The WP4 pixel DAQ project became possible *because* of the progress made in the trigger developments of WP3, software developments in response to the trigger needs (especially IPbus), and our long standing role in the Tracker, especially the readout, which was also being further developed by WP2. So removing WP2 and WP3 would abandon the activities for which we originally began the R&D and in which we took leadership roles and made major technical advances.

Conversely, removing WP4 would achieve insufficient savings, and would leave the new pixel detector without an off-detector readout system. The requirement to save a further £1.1M would halve the contribution to WP2, making most of the CBC ASIC development unviable and our role in future track-triggers in CMS, or would have a major impact on WP3, so limiting our role in the trigger, where we have a 1-2 year technical lead, have developed new concepts (IPbus and the Time Multiplexed Trigger) and have responsibilities for leadership, in the form of the Trigger Upgrade Project Manager, among other roles.

It is unlikely those remaining staff would stay motivated to continue the upgrade project. WP2 and WP3 are major planks of the Imperial College CMS programme and Imperial is the largest UK CMS group; the impact would be severe. Diminishing the activities at Imperial to achieve the savings would have a knock-on effect elsewhere, since the hardware has all been initiated there.

Similarly the other three UK groups would have most, or all, of their planned activities removed. All groups would lose the ability to further develop key skills and technologies; the dynamic and highly successful UK CMS detector-related activities would subside into something close to mediocrity; key staff with years of relevant experience will leave the project.

The loss of UK credibility would be immense and would have big implications for our future activities in CMS. Complete withdrawal from any WP would come out of the blue to CMS and we would not be able to point to any financial crisis as an excuse. The UK has a very high standing in CMS, both in honouring commitments and delivering challenging projects to the benefit of the experiment. There have been a number of occasions in the past where the UK stepped in to help solve very serious problems but these will be forgotten. In the future the UK would not be trusted as we are now and our future proposals will be questioned very critically.

There would be severe implications in the UK standing in CERN as well as CMS. The Phase I money matrix has been presented several times over the last two years, with indications of likely funding from all agencies. Gradually, CMS has managed to strengthen the commitments from most of the agencies, so that a substantial fraction of the Phase I funding is either in place or seems secure, on the basis of formal and informal assurances from agencies. Most recently the US made a very strong statement at the October 2012 LHC experiments Resources Review Board (RRB), "*affirming full proportionate support for the upgrades and intention to pay common funds in a timely manner that* *optimizes the success of the project* " [Quote from J. Incandela, CMS Spokesperson]. Since the US share of the Phase I cost is more than 30%, or around 23MCHF, this was very significant progress.

CMS is now in the process of setting up MoUs to put funding commitments firmly in place, for which STFC has pressed more than any other agency. The likely UK share of the Phase I upgrade costs has been included in all presentations and, while understood to be subject to peer review, its rough magnitude has not been questioned by STFC, who have been kept fully informed, nor the interest to play important roles in the pixel and trigger upgrades. At this point, therefore, it would be embarrassing to be seen to withdraw from one of those projects.

There would be long term implications for the future of the UK in CMS. By withdrawing contributions to either the pixel or trigger upgrade, we would lose the opportunity to make important contributions both now, in the next few years, and beyond that, simply on grounds of continuity, loss of expertise, and loss of leadership roles. This would certainly affect the UK participation in the subsequent upgrade activity and lose influence in important areas of the experiment, with knock-on implications for our physics and leadership roles.

Technically our activities would suffer, since we would give up the leading positions we have established on the basis of our demonstrated R&D progress and past contributions in the same areas to CMS. Whichever WP would be chosen for withdrawal, in all cases, there would be severe damage to UK credibility and a highly undesirable impact on CMS.

Withdrawal from	Contribution and expertise affected	Arguments FOR withdrawal	Arguments for NOT withdrawing
WP2	ASIC design and design of readout system. Track-trigger module development.	Latest item to be deployed in CMS, as it is aimed at Phase II upgrade, so CMS has time to minimise impact.	No other candidate group available to take over the work. Development time for complex ASIC and readout system is very long and time remaining is already tight. IPR, and practical, issues, mean R&D investment would be wasted, since it is very unlikely the present designs could be transferred to others. Expense of new ASIC development would have severe impact on CMS finances. Leadership role in this area, historically and currently significant, would be lost. Loss of expertise in track-trigger development which was originated by UK groups and is vital for the future CMS trigger and successful physics programme. Loss of expertise in novel assembly technologies and module development. Collaboration with CERN on advanced projects has ensured access to advanced technology, and close interaction with most advanced ASIC design expertise, currently 65nm CMOS.

WP3	L1 trigger and µTCA system development. Future trigger algorithms and LHC gtrigger implementation.	Other, US, groups *might* be able to take over the UK role in the project.	Immediate loss of leadership. UK staff member is present Trigger Upgrade Project Manager. Other roles would also be sacrificed. Technical leadership would be lost; UK is 1-2 years ahead of any comparable development. Opportunity to introduce original new trigger concepts would be lost, i.e. TMT. Loss of influence in trigger development, with important major implications for CMS and UK physics programme. Significant danger that CMS calorimeter trigger, with implications for physics programme, might not be completed.
WP4	Pixel readout and DAQ, µTCA system development. Signal processing for pixel data and direct experience of operating such a demanding sub-system in LHC.	µTCA development and maintenance of similar specialist skills can continue in the WP3, trigger, project.	No other candidate group is available to take over this work. Significant risk that off-detector readout system would not be completed in time, so late deployment of pixel detector, and physics impact. Loss of leadership in off-detector readout, with important implications for future readout systems and track-triggering. Loss of opportunity to gain experience in a strategic area of pixel development for future.

On the general point about excellence and leadership of the UK in the CMS upgrade activities, we draw attention to our submission to the 2013 STFC Programmatic Review process. A few points are repeated below:

The importance of upgrades to the CMS scientific programme was made clear in the CMS submission where we also explained the scientific importance of the physics under investigation in the near and farther future.

Much of the UK electronics infrastructure and skills, especially in advanced ASIC and high speed FPGA board design and use in systems, is unique and world class. It is also noteworthy that this partly resides in universities, and is complementary to the resources we use from RAL Technology. Other agencies have noticed our prominence in these important technical areas and have made efforts to emulate us. Italy, France and Germany, for example, at least in CMS, have little comparable strength in any of the electronic roles we have taken. They are all attempting to remedy this.

UK staff were early in contributing new concepts for the future Tracker and Trigger and were among the first to highlight the crucial importance of integrating Tracker data into a future CMS Level 1 Trigger to contain the overall trigger rate while maintaining efficiency for physics channels of interest. Important new concepts have emerged from our work, notably the concept of "stacked tracker" modules, which should be capable of suppressing low transverse momentum hits on-detector to ensure that data can be transferred to the trigger processors within the available bandwidth, and the idea of the Time Multiplexed Trigger, which provides boundary-free processing of calorimeter trigger data within a single FPGA, allowing new algorithms to be implemented which will be vital in the high pileup environment as the luminosity increases.

Other new developments have emerged from UK studies, such as the pioneering of μTCA telecoms hardware for LHC experiments, the hierarchical and expandable IPbus common firmware and software system to control and operate μTCA crates, and the design and construction of the first 130nm CMOS front end ASIC for readout of silicon microstrips in a future tracking detector. This work has been supplemented by simulation studies and software developments, which have been essential to motivate the hardware initiatives and define the Technical Design Reports now in preparation or under review.

The UK CMS groups had prominent leadership roles during CMS construction in all the areas of the project to which we contributed. Among these roles were the Deputy Spokesperson, later Spokesperson, from the inception of CMS until 2010. Additionally, the UK provided long term managers: Tracker Project Manager, Deputy Tracker Project Manager, Upgrades Project Manager, Tracker Upgrade Coordinator, Endcap ECAL Project Manager, Global Calorimeter Project Manager, Trigger Upgrade Project Manager, Trigger-DAQ Institution Board chair, among others, supplemented by representative roles in sub-detector management.

During planning of CMS upgrades, UK staff have continued to be very prominent in steering the upgrade work, particularly in the areas which involve the most substantial changes to the detector and require novelty and innovation, which are the Tracker and Trigger.

Having built up significant roles in the upgrade activities, which are quite closely related to the physics interests we have in CMS, any undermining or sacrifice of our positions would certainly represent a loss to the UK. Future opportunities to take leadership positions in the experiment will arise from proven records of project delivery, and the experience of taking responsibility for leading edge projects is important for career development of younger staff. The scientific rewards will be significant, and the upgrade developments will maintain UK physicists and engineers at the forefront of technical innovations.

2 Costs of the Phase II upgrade of CMS. Can you be more precise about expectations, especially for likely staff costs?

At this stage it is quite difficult to provide improved cost estimates, which depend on the long term scientific goals of CMS, the required duration of future operation and the longevity of the existing detector, as well as potential performance gains which might be achieved.

It is worth recalling that the major drivers for the overall upgrade programme are the needs to replace the Tracker because of radiation damage after a decade of operation, and to upgrade the Trigger to cope with the highly challenging pile-up conditions. Operation of all sub-detectors becomes even more challenging in the higher pile-up conditions of 50ns bunch spacing, which is not ruled out because of concerns about electron cloud effects and vacuum quality in the LHC using 25 ns bunch spacing. These were phenomena which were not anticipated during the LHC design whose impact was only appreciated shortly before operation began.

It is not likely that the Tracker will suffer less radiation damage during operation, so its replacement is inevitable. The discovery of the 125 GeV boson has increased pressure to maintain performance with low transverse momentum and energy thresholds in the Level 1 trigger simply because a large fraction of the boson decay products will have low momenta, just from kinematics. The study of the Higgs branching ratios and couplings will inevitably form a major part of the long term LHC programme, since their precise evaluation is crucial to understanding ElectroWeak Symmetry Breaking in detail and the alternative to making these measurements at the LHC would be an electron-positron collider, which would be significantly more expensive and unlikely to be available on the same timescale. In addition, it is still widely believed that further evidence of new physics will be found, and the low mass of the Higgs boson adds support to ideas connecting this to SuperSymmetric models.

In view of the requirement to extend the operating life of CMS up to and even beyond 2030, questions must be asked about the longevity of the detector and the possibility for previously unforeseen maintenance and replacement of vulnerable components. These are thought to be most likely among electronic boards; although not the ASICs and other similar components. Replacement, for example, of the ECAL Front End boards would also permit increase of the latency and trigger rate which are limited by the design of these electronics. While this would be desirable, the cost and impact of doing it, which requires moving services on the magnet vacuum tank, is a major undertaking. Changing the latency and L1 rate means all CMS electronics must be adapted to the new requirements.

All these points are under discussion and decisions are expected during 2013. Until then the best estimates of future costs remain those in the CMS Expression of Interest dating from 2007, which foresaw an overall upgrade (materials only) cost of 175MCHF, excluding R&D. Much has changed since that time, including the requirement for separate Phase I and Phase II upgrades. However, some of the cost estimates still look reasonable; for example the pixel system was expected to cost up to 30MCHF, while the 4 layer Phase I replacement system will cost 17MCHF. Similarly the new trigger was estimated to cost 20MCHF; the Phase I upgrade is expected to cost 5-10MCHF. A new Outer Tracker was estimated to cost 90MCHF. However only 15MCHF was allocated for muon systems and calorimeter upgrades and 15MCHF for infrastructure; these may be underestimates.

Perhaps a realistic estimate of the Phase II upgrade cost would therefore be 200-300MCHF, in today's currency, including the remaining R&D required. In that case the UK share would be 8-12MCHF, over a period of about 6-7 years, anticipating possible slippage in the LS3 date. No judgement is possible of the sterling value, except to note that in 2007 the exchange rate stood at about 2.3 CHF/£ while today it is about 1.45 CHF/£.

Staff costs are harder to estimate. Probably the most realistic method would be to assume that funding constraints would not permit significant increases in annual spending compared to the present. In that case, given that the activities expected from UK groups would build on the present planning and be in the same areas of Tracker and L1 Trigger, it would be reasonable to assume a request similar to the present one, namely £7-8M in new staff requests, with £5-6M in existing staff, and a travel request of ~£0.5M, leading to a total of £12.5-14.5M. Working allowances and materials costs not accountable in the CMS cost book might add a further £1.5M to the request.

In today's currency, this leads to a total expected to be spent over about 6-7 years:

Share of CMS construction:	£5.5-8M
Staff costs (new and CG-funded):	£12-14M
Travel, working allowance, materials not in CMS construction budget:	£2M