9 September 2012

CMS Upgrade: Responses to referee comments

We found the responses from the referees to be overall extremely positive and we are grateful to them for their insight and support. The referees noted positively the importance of the proposal as well as its ambition, technical excellence, the UK track record of innovation and leadership roles and our management capabilities. We were pleased by the strong endorsement by one referee of the risk management strategy we have adopted. Obviously we cannot claim credit for the physics background but, since the proposal was submitted, the hints and inside knowledge of a 125 GeV mass boson discovery by CMS and ATLAS have been confirmed and provide further motivation for the need for the detector upgrades. The referees noted the crucial importance of the areas where we focus our efforts, namely advanced tracker readout and the trigger. In addition to exploring the properties of the newly discovered particle, CMS is increasing its efforts to uncover SUSY or other evidence of new physics. UK groups are very active in Higgs, SUSY and heavy vector boson searches, as explained in the CMS grant request of February this year, which are the most likely to uncover new physics, along with complementary exotic activities such as searches for long lived particle decays. All of these rely heavily on the exceedingly good performance of tracking and Level-1 triggering in CMS.

One referee (YGYCX) commented that the programme seems to "occupy an R&D niche within CMS with minimal competition from international collaborators". We doubt that this was intended as a disparaging comment but, as it might be misconstrued, would like to respond. CMS, while a strong collaboration, does not have enough resources to permit duplication of activities. "Shoot-outs" have been very rare occurrences in CMS history and shared activities are normal, with little or no unnecessary overlap. The UK has established its footprint in CMS by demonstrated success in the Tracker readout, Trigger and DAQ activities, and has also taken a lead in upgrade R&D in its areas of expertise. We did not stress our past successes during CMS construction in the proposal but would be happy to expand on our those achievements on which this project builds, if desired. The pioneering 0.25μm CMOS ASIC work is one example and the development of FPGA-based systems for both Tracker and Trigger is another. We have continued to maintain and operate our systems, as CMS expects, and our grant proposals have explained our contributions and how they complement our physics analysis studies.

We will attempt to address the criticisms, most of which are probably a consequence of the strict page limit for the proposal and the consequent choices which have to be made between explaining sufficiently the motivation and background to the project, or justifying the details.

Referee YGYCV

Academic leadership in WP4

We find the comment that there is a lack of academic leadership a little surprising, especially following the comments about the PI. In fact, the PI has had a long standing role in the management of the Tracker and this continues today with new pixel management tasks supplementing the operational management of the microstrip tracker which is now in a steady state, but still requires attention. The PI was until last year leading the management of the Tracker Upgrade. He was responsible for proposing the development of a pixel Tracker FED as a UK project, since it offered an ideal opportunity for the UK to take a significant role in the pixel system, in which we have not previously played any part.

Although the Tracker comprises two major sub-systems- pixels and microstrips, historically these have developed almost autonomously within the Tracker collaboration under a single overall Project Manager. There has been a constant effort to avoid unnecessary duplication; for example the pixel tracker uses optical links, Front End Controllers, cooling and power systems which originate in the microstrip tracker, and also followed the Tracker on the same ASIC technology route. However, in most cases these components are customised specifically for the pixel system, so the two communities work closely but with little direct overlap. This is now changing as the benefits of a homogeneous DAQ are evident, and groups from the microstrip tracker with, e.g., sensor or module construction expertise enter the pixel system to contribute to the upgrade, both to learn about pixel technology and to ensure the new detector can be safely delivered in time.

Thus the UK is entering a new but related area and it was natural to exploit the expertise we already have to contribute to the pixel system while learning about its special requirements. Our ASIC expertise is not needed (PSI is very competent) and we have commitments to WP2 under way for the longer term, but our expertise from the microstrip FED, online DAQ, and unique knowledge of μTCA technology, will significantly strengthen an area of the pixel system which was previously weak. Strategically, it also places the UK in a strong position to continue this role into later stages of the upgrades, when the whole tracker is completely replaced. It was also natural to exploit achievements in the R&D activities so far, where Bristol and Imperial developed the new IPbus control system for μTCA and the μTCA hardware and firmware. More recently RAL PPD has developed a test board for the new pixel front end chip readout which will help guide the new FED specification, as well as build on past contributions to Tracker work in offline software. Academics in all three institutes are responsible for their teams, and we have identified a key role for the Brunel group which also builds on their real past contributions to the Tracker Data Quality Management and software development.

Regarding the WP4 project managers, if this was what the referee had in mind, we have two young but highly competent physicists who should now acquire experience in managing a significant size project (a relatively rare opportunity) and who should get the credit for doing so. We cannot easily anticipate career progression, but it is crucial to hand over such responsibilities to younger staff.

Finally, as the construction project evolves, we expect other academics to take an increasing interest in the tracking system. The Tracker and ECAL are the most crucial detectors for the UK physics programme and a deep understanding of the data has been important to our, and CMS's, progress.

Role of Brunel

We note the comment of referee YGYCW that the Brunel group "knows how to focus in their areas of expertise". The main responsibility of the Brunel group in WP4, building on their expertise in contributing to both commissioning of the existing strip tracker and to the current upgrade R&D, will be integrating and adapting CMS data quality monitoring software for the upgraded pixel detector, which is intimately tied to the FED online software. Significant experience exists at Brunel of optimizing memory usage and speed of reconstruction which will be of considerable value here. A second important contribution from Brunel, in collaboration with Strasbourg, will be the database interfaces and integration needed for FED (and FEC) configuration. Support for integration with offline software, based on the significant prior experience of Cole (contributions to strip tracker DAQ software and commissioning), Reid (optimization of CMS tracking offline reconstruction software and contributions to phase-1 upgrade design and software efficiency for the new tracker) and Hobson (software project management experience and, with Reid, development and evaluation of new algorithms for DQM), will be provided. A recent group member (Powell) is a systems engineer with extensive experience in designing software for on-line control of large-scale systems. He has very considerable experience in database design and optimization and has contributed to the MICE project on software development, release strategies and new techniques for fast track reconstruction.

Could other groups deliver other (non-μTCA) parts of pixel DAQ?

In short, there are no significant other parts of the pixel DAQ than the μTCA system. The costs in our proposal are only for the main components, i.e. the FEDs, ancillary boards (MCH and AMC13) and crates. The optical receivers and fibre links are provided to us and funded separately by the pixel project. Strasbourg will contribute to the DAQ implementation by taking responsibility particularly for the Front End Controller (FEC) implementation, which builds on their contributions to the present tracker. However, they are not able to take responsibility for hardware delivery, as French contributions had already been committed to the pixel cooling system.

New RA posts at Bristol, Imperial and RAL PPD

The work plans which have been devised for each WP were used to estimate the staff effort required in some detail and thus motivated the new staff requests. Details of the roles of new RAs were mentioned by all referees and further details are given later. The main requirements for each post are listed below:

Bristol RA 1: WP3

The RA will undertake two key tasks. Firstly, develop the automated firmware build and test system which is required to allow TMT firmware components to be integrated. The quantity and complexity of 'algorithm' firmware required for the TMT is extremely substantial compared to previous systems using much smaller FPGAs. A software-style build system capable of integrating with a multi-contributor firmware repository, and carrying out automated regression tests and comparison with the offline trigger emulator, is essential to address this issue. Secondly, the RA will take responsibility for development of the emulator modules that correspond to running trigger algorithms, and allow real-time and offline checks of trigger operation, and will provide tools to allow automated data capture and check during both commissioning and operation phases. We expect the RA to play a central role in trigger installation, commissioning and first operation, and to be on LTA at CERN for an extended period during these phases.

Bristol RA 2: WP4

This RA will take the central role in the development of the low level control software for the pixel FED. This software will interface to the higher level CMS control system to allow configuration, control and monitoring of the entire pixel readout system. This is a substantial task, which later in the project will involve a great deal of work with the final pixel readout system. The RA will work closely with WP3 online software counterparts at Bristol and elsewhere, to exploit common components and approaches wherever possible. The RA will play an important role in system integration at CERN, and will also take responsibility for ensuring system reliability under high luminosity and heavy ion running conditions.

Mark Grimes:

We were asked to distinguish Mark's role from that of the Bristol RA 2. Whereas the new RA post is focused on the online systems, Mark has significant experience on offline software development and testing for the upgraded pixel system, working at all levels from detector enumeration and raw data handling, to evaluation of physics performance. He will exploit this knowledge as the emphasis turns from simulation studies to the interfacing of the new data formats to the CMS offline software, working in conjunction with existing CMS collaborators at FNAL and elsewhere to develop and test the relevant modules for unpacking, filtering, reconstruction and DQM. In this task, he will also liaise closely with the Brunel effort for DQM, providing technical expertise where required.

Imperial RA 1: WP2 (60%) & WP4 (40%)

This RA will complement and work closely with M Pesaresi, J Fulcher and M Raymond and contribute to two WPs. The evaluation work and module studies in lab and test beam using the

CBC require a steady 1 FTE in addition to Mark Raymond and a project student (summarized on p14 of the proposal); this is therefore foreseen to be shared by M Pesaresi and this post.

For WP4, the Imperial 1 FTE total, again contributed by M Pesaresi and this post, is required to contribute firmware and online software with particular emphasis on the DAQ interface, and the local DAQ software that interfaces to it. The RA will apply the DAQ firmware and software expertise, working closely with J Fulcher, to support commissioning of the detector and evaluation of its performance.

We feel that it is desirable to contribute work on both WPs to gain experience and to permit flexibility in the interests of the project. There will be a lot of overlap between the WP2 and WP4 tasks, since they will share the same hardware and most of the firmware. The local WP4 DAQ system is very closely related, if not identical, to the future test beam DAQ system for WP2.

Imperial RA 2: WP3

The RA will support G Iles and A Rose working on algorithm simulations, firmware and hardware commissioning. The MP7 and the system will require extensive lab evaluation and algorithm development, including detailed quantitative performance evaluation, as well as eventual commissioning. The profile for the post is likely to be a physics background in the trigger, with existing, or capability to develop, firmware and online software skills. We expect the postholder to spend significant periods in CERN to contribute to the parallel operation, debugging and evaluation of the new trigger, and possibly be on LTA during the commissioning of the trigger. Imperial has nearly all the hardware responsibilities for WP3 and WP4. The load on those who have those tasks must not become excessive; hence the requests we have made.

Imperial RA 3: WP3 3 year post, from FY16/17

The trigger is a sub-system which can only finally be commissioned in the experiment with real data. It is crucial that this is fast and as efficient as possible with key individuals being available on the spot, and able to react immediately to identify bugs and correct firmware, then ascertain the results are in full agreement with emulations at bit level. It is even more critical when the experiment is undertaking an upgrade and most other systems will be fully ready to go and therefore impatiently waiting for the trigger.

Experience from commissioning the GCT and our estimates of staff available for this project, particularly in CERN, motivate this post to focus on the trigger commissioning. G Iles will continue to have operational duties maintaining the existing trigger while that is in operation and A Rose will be based at Imperial responsible for the test system where firmware is developed, evaluated and debugged. The post holder would spend significant periods in CERN, probably on LTA during the final commissioning of the trigger.

RAL PPD RA 1: WP3

This RA will carry out two tasks. Firstly, participate in development of a new coherent offline software framework for the L1 trigger. This encompasses trigger data handling code within the CMS offline and HLT software packages, the trigger emulator framework allowing simulation of trigger performance, and tools for trigger data quality management and performance monitoring. The replacement of the calorimeter trigger, along with parallel developments in the global and muon trigger components, require substantial changes and improvements to the offline software framework, and CMS will take advantage of the LS1 shutdown to redefine the architecture of the system to improve its future maintainability, and introduce common approaches across subsystems. The RA will play a key role in the confirmation of the correct performance of the parallel trigger system during extended operation, an essential criterion in switching to the new L1 system. Secondly, they will contribute to the programme of simulation and design studies for the post-LS2 track trigger, in direct support of UK academics involved in this area. This will focus mainly on the design and emulation of the off-detector track finder component of the system, corresponding to the UK hardware and analysis interests, and complementing the WP2 work on the readout electronics. In the latter part of the project, the RA will play a key role in the validation of the track trigger prototype through detailed emulation.

RAL PPD RA 2: WP4

Required to oversee full system tests, certifying and improving firmware and online software initially for prototypes at RAL, and later for larger test systems at CERN (TIF/P5). Responsible for integration with the CMS trigger system. Later on, participation in commissioning and performance evaluation of the upgrade detector.

Justification of project students

Space did not permit us to justify individually each request, but we did believe we had conveyed the purpose and importance of these studentships:

P11, WP1 (Management): Five project studentships are requested. We regard them as essential; excellent training opportunities are provided, with highly transferable skills gained; in addition it ensures throughput of trained, skilled candidates for posts during the subsequent operational period.

P14, WP2: The success of the CMS chip development has resulted from close collaboration between microelectronics design engineers at RAL and university staff and students, who perform the detailed studies on chip and module performance.

P19, WP3: The current R&D project has provided opportunities to train several students in relevant areas, providing a pool of effort. We also propose to engage project-funded postgraduate students at Imperial and Bristol in the work package. The design and delivery of a trigger system, from conceptual design, through construction, to final commissioning represents a truly excellent training opportunity, which is heavily interlinked with the CMS physics programme.

The current R&D project has directly funded only one project student but during the development of the CMS detector the work of the PI at Imperial was supported by a series of excellent students, some of whom were recruited as CASE students in collaboration with RAL TD. Many of them have gone on to notable positions in the field, including appointments in CERN and CMS. Examples are Karl Gill (CERN staff, current CMS Tracker PM), Greg Iles (CERN Fellow and now Imperial engineer), Jan Troska (CERN staff post, optical links), Jonathan Fulcher (industry, and again Imperial DAQ expert), Matt Noy (CERN Fellow on NA62, then staff post) John Jones (Princeton, and his own company), Mark Pesaresi (RA at Imperial College), Andy Rose (Imperial RA, now Imperial staff engineer), among many others.

The project represents a superb training opportunity for technically oriented students. Good students specialising in these areas learn firmware and online software skills fast and become very effective contributors to our work. They are crucial because these specialist skills are not easy to learn, being quite unlike data analysis software, but are vital to deliver the working hardware fully implemented in the experiment, about which they are well informed, motivated and able to innovate. Being young and not yet established, they are also mobile and able to relocate if needed. In contrast, RAL engineers do very solid work but are unable or unwilling to work for lengthy periods in CERN, e.g. during commissioning activities, and they are costly to relocate. They are usually less flexible, which is essential in a fast moving environment, such as the CMS trigger or online DAQ development. Staff turnover is both desirable and to be expected during the project lifetime and we will need to recruit replacement. Suitable young candidates are not easy to find, except from similar projects in our field, which are relatively few.

We have noted some of the tasks to which students would contribute: chip and module evaluation in labs and test beams, development of online DAQ systems, trigger algorithm firmware and software development, performance evaluation, including simulation studies. All of these are crucial, and it was not an oversight that we imagined that a student might work on more than one WP during their training, because as well as the experience gained, flexibility and adaptability are important aspects of the skills to be gained.

Details of equipment costs

We regret that our estimates were not as clear to the referees as we had hoped. Page 10 "Resource Estimates" of the proposal gave an overview of the process we have followed. WP2 details are in the proposal text and seem to satisfy the referees, being mainly ASIC manufacture costs. The tables for WP3 and WP4 summarise estimates from the much more detailed Excel workbooks into the main headings. WP3 and WP4 cost breakdowns are quite similar, except for details of the components, which are mainly FPGAs, whose requirements are different since WP4 is able to target a less expensive part with lower speed interfaces.

The main subheadings are Preproduction and Production costs of the trigger or pixel FED boards, and Working Allowance, which is explained on page 11 of the proposal. In addition, for WP3 there are Infrastructure costs, explained on p19 of the proposal, which is a pro-rata contribution to the optical links and transceivers needed to build the entire trigger system and which is therefore a shared cost not associated with a single agency or major trigger component. WP3 also requests sub-contract costs for Iceberg Technologies which is further justified later in this document.

Both WP3 and WP4 request modest amounts of local institute support costs which were detailed in our earlier answers to the STFC office.

As the tables and text indicate, the major costs of both WP3 and WP4 are in the board production costs which inevitably lead to substantial totals, but which are quite in line with the expected UK share of the CMS Phase I upgrade cost.

For WP3, the required cards, particularly the MP7, are very high performance (~1Tb/s) so the unit cost is high, mainly driven by the high specification FPGA. The trigger board costs approximately $15.6k (actually there are two processor types, one with fewer optical receivers costing $14.1k). This is made up of PCB manufacture and assembly $1.8k, FPGA $6.8k, components $1.9k, optical transmitters and receivers $5.2k. There are two types of trigger crate, using the two processor board types, costing $205k and $154k. The crate costs include processor cards, ancillary cards, power supplies, backplane and mechanics, and cables. Two crates of each type are needed, with one spare of each type. The spares will be used as test beds in CERN. With pre-production and minor costs for main power and interfaces, this arrives at the £774k quoted in the proposal assuming a board production yield of 85%.

WP4 costs are estimated in a broadly similar way using a modified workbook. The requirements are for a high density, low bandwidth solution, so the FED cost can be reduced compared to the MP7. Fewer high speed serial links are required and therefore a less expensive FPGA part can be used. Optical components are either provided to us as part of the pixel project, or are relatively inexpensive or few.

For WP4, the FED cost is approximately $8.3k, made up of PCB manufacture and assembly $2.3k (including mezzanines), FPGA $4.0k, components and SFP transceivers $2.0k. Five FED crates with a unit crate cost of $98k (including the cost of ancillary cards, power supplies and backplane) are needed; plus one fully populated spare, which will also be used for DAQ development. Prototyping costs total £114k (unlike the MP7, the FED prototype does not yet exist) and will be a necessity for the early integration and commissioning of the DAQ with the pilot and phase I pixel systems. With pre-production boards and the cost of the required crate PCs, interface cards, cables and main power supplies, this arrives at the £651k quoted in the proposal, assuming a board production yield of 85%.

Total costs and possibility of scaling down

Obviously we believe that the equipment costs for WP3 and WP4 represent an accurate estimate of the two systems we propose to provide to CMS. This will be needed, even if UK funding is not granted in full for the material requests. Although it is not intended as an argument, it should be pointed out that the materials costs of the two projects are quite close to the fractional share expected from the UK to the total Phase I CMS upgrade budget, which is around 2.6MCHF, i.e. 4% of the total upgrade cost whose latest estimate is 66MCHF. Not all the costs are eligible to be considered as cost book items, e.g. local institute expenses and some prototyping.

Contribution of other CMS groups contribute to chip production costs

When production starts we do expect the costs to be shared with other CMS groups. Other than CERN, who are contributing substantially to module development, we do not yet know which groups will share the costs. Other groups currently involved in this part of the upgrade project are Lyon and Strasbourg but it is too soon to assume this will be a French contribution.

Referee YGYCW

Effort required for WP3

The WBS for WP3, shown in the proposal, was used to estimate the required effort. The current level of effort (ignoring new posts) assigned to WP3 to "Hardware / firmware systems, WBS" is 1.3 FTE, of which 0.25 is for layout support. The remaining 1.05 FTE is insufficient to develop; firmware of a new system; procure and validate cards; assemble a dedicated system for integration and testing; and assemble the final system. This therefore needs to be strengthened by at least 1 FTE, which will probably need to be on LTA at CERN. Procurement of large projects is complex and therefore 0.17 FTE was assigned to RAL who have significant experience.

The "Online software" section of WP3 WBS has 1 extra post to reach 1.9 FTE in total. The extra post would be filled with R. Frazier who has provided excellent software support to the current GCT. 1.9 FTE is not excessive for a project of this scale and will probably need further strengthening during critical periods from WBS section "Trigger studies + offline software".

At present the "Trigger studies + offline software" section is relatively small (0.55 FTE) and it is spread over 3 personnel. Strengthening provides much better guidance to the firmware & hardware algorithm developments. This section has traditionally performed other roles as well (e.g. DAQ unpacker software). It also validates the final system by providing bit level hardware emulation. In this phase the section is extremely valuable in helping commission the system and running stress tests. Given the extremely low FTE currently assigned it makes sense to strengthen this section by 2 FTE.

Over the R&D phase roles in the μTCA and μHAL developments have changed depending on requirements and skills. For example, Andrew Rose initially worked on IPbus firmware, then developed the C++ interface and has now returned to system development (mix of firmware / software).

Something similar will occur during the project phase (i.e. initially there will be extensive hardware testing, followed by a shift to firmware development). Software development will continue throughout the project, but as the project draws to an end the bulk of the personnel will be used to provide software (e.g automating processes, diagnostics, etc).

Distribution of WP3 tasks between institutes and management structure of WP3

The two WP3 managers are deeply involved in the CMS trigger project and provide, with the PI, the top layer of local management. Below that there are several major tasks to be covered which we plan to be overseen as in the table below. We do not foresee problems of sharing responsibilities or communicating; except for RAL the groups and individuals have been working together for a long time, and very successfully, as the development of IPbus and μHAL attests.

The upgrade effort at RAL will be bolstered during the years 2012-14 by additional academic effort from D. Newbold, who will undertake a 50% secondment to the lab from Bristol. This will ensure RAL's speedy and effective entry into the L1 trigger project, and help cement cooperation in firmware and online software between WP3 and WP4.

Area	**Manager**	**Contributors**
Hardware and procurement	G Iles	Imperial College
Firmware	G Iles	Imperial and Bristol
Online software	R Frazier	Bristol, Imperial, RAL
Offline software	J Brooke	Bristol, Imperial, RAL
Algorithm development	D Newbold	Bristol and Imperial
Commissioning	A Tapper	Imperial, Bristol, RAL

Bristol past delivery record

It is true that Bristol ran into problems with the delivery of the original GCT project. However, this was partly due to the fact that the electrical link technologies which they were forced to deploy for the project to provide connections to other parts of the trigger were running beyond their stated specifications. The cables carrying those signals were specified by other members of the trigger also gave rise to serious problems. Later it was obvious, with the benefit of hindsight, that the original project was technically much more demanding than originally planned and the engineering support requested was insufficient.

Imperial joined the project, with engineering support partly from CERN (mainly G. Iles), and profited from the evolution in FPGA technology to deliver a more compact system with fewer boards. However, Bristol did still play a significant role in the new GCT; they provided much of online software, from which this project also profits, and gained new expertise. They have also contributed notably to the operation of the GCT and trigger simulations. In the upgrade R&D project they have delivered impeccably, and share significant credit for the IPbus system.

Sharing of online software tasks

See earlier comments.

Referee YGYCX

R&D "niche"

See comments in preliminary remarks.

Justification for contributing to WP4

See response to Referee YGYCV.

WP4 effort requirements

The WBS is provided in the proposal and has effort estimates. There is another level below where resources were matched to tasks.

WP4 can be characterised as three distinct phases (p23) covering initial prototyping and testing (WBS 4.1, 4.2, 4.3, 4.4); final system production and commissioning (WBS 4.5, 4.6); and eventually DAQ operation and detector studies (WBS 4.7). The hardware effort required to complete the first two phases (covering four years) will be 5.2 FTE, including a level of support for board rework and the production of custom mezzanines. A continuing low level of hardware support at 2.2 FTE is required in the last phase to cover repairs to boards and mezzanines during operation.

A significant level of firmware and online software effort will be necessary in the first phase of the project. This includes implementation of the base board firmware and testing procedures, TBM/ROC decoding and multi-channel synchronisation, DAQ link prototyping and back end firmware, porting and integration of FEC firmware, as well as the design of an online software framework. The requirement would be 12 FTE over the duration of this phase (covering two years). This must continue into the second phase (covering the following two years) where the emphasis will be on supporting the integration effort (between the final boards, DAQ and trigger control systems) and commissioning and algorithm development (with the pilot and Phase I detectors) when much of the work in the first years will converge. A limited amount of firmware and online software support will be required in the final phase of the project at the level of 3.75 FTE over two years.

While the testing and integration effort will be supported by engineers assigned to firmware and online software development, this will work will require 7.4 FTE over the first two phases of the project (covering four years), especially when the focus of the testing moves to CERN. An effort on offline software development, including DQM tools and analysis, is also necessary and will become a major component as the project moves from the construction phase to the commissioning and operations phase. To fully exploit the work undertaken in the initial phases, much effort will be required to commission and measure the performance of the Phase I detector and DAQ during collisions, contribute to tracking, hit multiplicity and luminosity studies and to prepare for Phase I operation at the level of 13 FTE over two years.

Possible WP4 scaling down to know-how only

As described already, we believe WP4 is an excellent project, both for the pixel upgrade and strategically, and the materials cost is within the range of expense which the UK should, according to agreed cost sharing algorithms, be able to deliver to CMS. The pixel project cost of 17.1MCHF has contributions from all major agencies broadly in line with such expectations (as the PhD author algorithm can only be applied at CMS level, not sub-detector level). So CMS collaborators are "stepping forward" to bear equitable shares of the project cost. In addition, it is probably not practical to provide know-how only, which would even in the best represent only a short term contribution while others learned the technologies involved. Diffusion of knowledge is commonplace and beneficial in particle physics but we would hope not just to pass on our experience alone to others, and not reap the benefits of our substantial investment of intellectual input, effort and UK-provided R&D funds.

Function of new RA posts and detailed justification

See response to Referee YGYCV.

Cost of Iceberg Technologies

The founder of Iceberg Technologies is John Jones, who is one of the originators of the Time Multiplexed Trigger concept. He is an astonishingly talented individual with remarkable engineering (and physics) skills and he made major contributions to the Global Calorimeter Trigger, and other projects, while still a PhD student. Subsequently he moved to Princeton, on a special Fellowship still working on CMS and formed his company working on projects for such sponsors as LBNL, Oak Ridge and the Elettra synchrotron in Trieste. He has remained in regular contact and has made formal and informal contributions to the MP7 design. He is an exceptional firmware developer.

We would like to retain his contributions to the project but he is not available to be hired as an RA for probably obvious reasons. He is ready to contribute to specific demanding tasks, as he retains a deep interest in the project. Up to now these have included the layout of crucial parts of the MP7, such as the 10 Gbps interfaces, and certain firmware modules. The funding would also allow us, if necessary, to call on his skills should we encounter problems or, more likely, anticipate them using his extensive experience. Some of his expertise is also effectively transferred to us so younger staff will benefit. We consider him excellent value for money.