December 15, 2006

US CMS Heavy Ion Physics Program Personnel and Funding Resources Input to NSAC Long Range Plan

M. Ballintijn[7], O. Barannikova[2], R.R. Betts[2], B. E.Bonner[9], W. Busza[7], M. Calderón de la Barca Sánchez[1], D. Cebra[1], P. Constantin[5], G. Eppley[9], E. García[2], N. George[3], O. Grachov[4], S.V. Greene[10], D. Hofman[2], R. Hollis[2], A. Iordanova[2], J. Kapusta[8], G. J. Kunde[5], C. Loizides[7], W. Llope[9], C.F. Maguire[10], A.C. Mignerey[6], C. Mironov[5], M. Murray[4], G. Mutchler[9], E. Norbeck[3], Y. Onel[3], C. Roland[7], G. Roland[7], S.J. Sanders[4], G.S.F Stephans[7], M.B. Tonjes[6], J. Velkovska[10], R. Vogt[1], B. Wyslouch[7], P. Yepes[9]

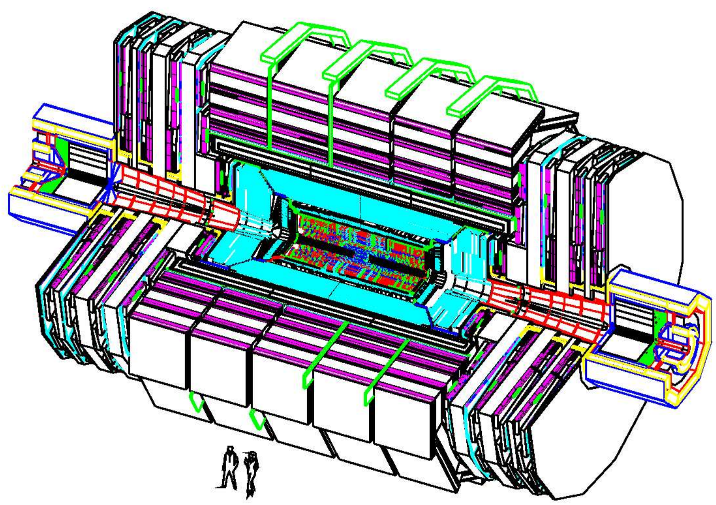

[1] University of California at Davis, Department of Physics, Davis, California 95616

[2] University of Illinois at Chicago, Department of Physics, Chicago, Illinois 60607

[3] University of Iowa, Physics and Astronomy Department, Iowa City, Iowa 52242

[4] University of Kansas, Department of Physics and Astronomy, Lawrence, Kansas 66045

[5] Los Alamos National Laboratory, Los Alamos, New Mexico, 87545

[6] University of Maryland, Chemistry Department, College Station, Maryland 20742

[7] Massachusetts Institute of Technology, Laboratory for Nuclear Science, Cambridge, Massachusetts 02139

[8] University of Minnesota, School of Physics and Astronomy, Minneapolis, Minnesota 55455

[9] Rice University, Bonner Nuclear Laboratory MS315, Houston, Texas 77005

[10] Vanderbilt University, Department of Physics and Astronomy, Nashville, Tennessee 37235

1 Introduction

This document summarizes the personnel and funding resources pproposed for the US portion of the CMS Heavy Ion physics program. Heavy ion physics is an integral part of the CMS experiment, which will be operated during both the p+p and Pb+Pb beams at the LHC. Physicists in the heavy ion groups are full members of the CMS collaboration and actively participate in the detector construction as well as data analysis for both p+p and Pb+Pb collisions. One specific responsibility of the heavy ion members is preparing the experiment for Pb+Pb data, in particular the trigger and online software. Other more general responsibilities include experiment maintenance, shift taking, data processing, and data analysis. Bolek Wyslouch of the US CMS heavy ion group has been appointed coordinator of the Heavy Ion subgroup which consists of physicists from around the globe with the US group playing a leading role. The prominence of US physicists in this area is illustrated in Fig. 1.

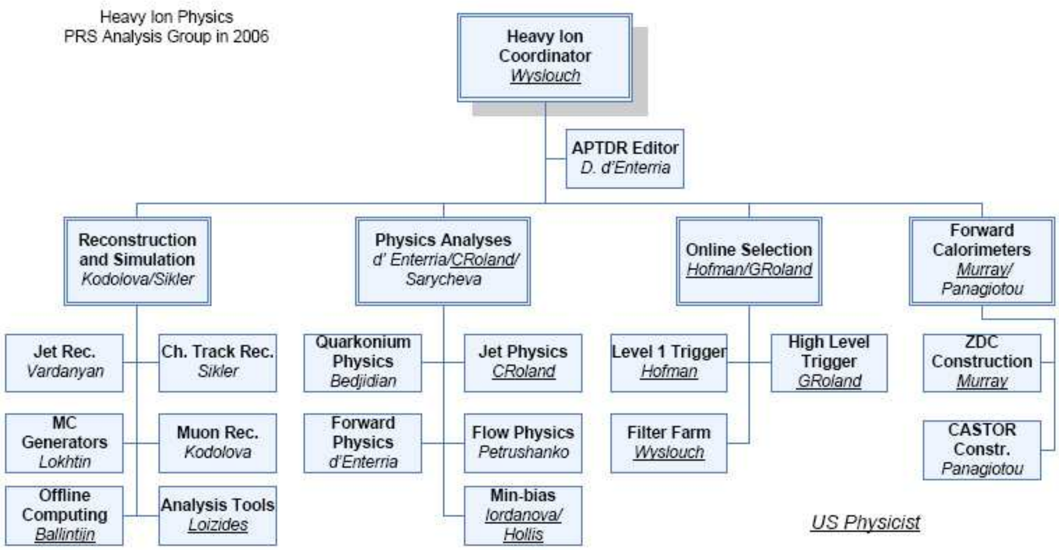

Figure 1: Organization of the heavy ion physics group in CMS. To illustrate their importance in the management structure, the names of US participants are underlined.

In addition, heavy ion physicists are involved in several areas of hardware construction and preparation. In particular, the US groups are building the ZDC calorimeter and part of the DAQ Filter Farm. The hardware and operating costs of US heavy ion participation can be separated into five categories (see Sections 2, 3, and 4 for detailed discussion) :

- Experimental equipment to be installed at CERN (ZDC and DAQ Filter Farm, see Section 2);
- Yearly contributions to the running costs of the experiment which are proportional to the number of participating PhD level physicists (Category A, see Section 3);
- Maintenance and repair of the installed equipment (Category B, see Section 3);
- Offline computing resources for data processing and data analysis (see Section 4);
- Increase in personnel and travel cost for on-site participation at CERN above what was typical for the RHIC experiments (see Section 3).

2 Capital Funding Request

Extensive development of the High Level Trigger is necessary in order for CMS to fully exploit its potential for heavy ion physics. Therefore, the US CMS heavy ion group will focus its effort and resources on the development of the HLT and the software and hardware infrastructure of the Filter Farms. In addition, we have proposed to contribute to the experiment by purchasing, testing, installing, and commissioning parts of the trigger in time for the first heavy ion collisions at the LHC.

The development of sophisticated real-time algorithms requires long, detailed, and careful testing and debugging. Hardware purchased relatively early in the project will be used to establish a test bed for various aspects of the software and data handling. We plan to purchase the computers in the US and ship them directly to CERN, where they will be extensively tested using standard CMS "burn-in" procedures. We will use all available funding to buy DAQ PCs with the exception of the initial purchase of machines which will be used for development and testing of HLT code.

The initial significant block of CPUs will be in place for the first Pb+Pb run which is projected to be at about 1/20 of the full design luminosity. The processing power of the partially completed HLT will be needed to ensure the cleanest possible event sample and also for data compression to maximize the event rate to tape. The second major purchase is scheduled to allow some contingency in the plan for installation and testing in time for the first full luminosity heavy ion beams. In order to maximize the CPU power of the full system, the final large purchase will take place as late as possible before this second run.

The construction of the development system will require technical help as well as hardware including racks, power strips, and network switches. We estimate that we will need 0.25 FTE technician in FY2007 to set up the test station and conduct hardware checks. A breakdown of the DAQ-related budget items in each of the fiscal years is shown in Table 1. The additional hardware needed to set up the test facility is included in the "Equip." column. As equipment, the hardware and PCs will not incur overhead charges. The technician cost includes benefits. The last column lists expected overhead and other institutional charges.

Year	PC Units	PC Cost	Equip.	Tech. (SWEB)	Overhead	Total
		Costs in k$				
FY2007	30	93	6.5	21.6	18.9	140
FY2008	200	620	0	0	0	620
FY2009	200	620	0	0	0	620
FY2010	200	620	0	0	0	620
Total	630	1953	6.5	21.6	18.9	2000

Table 1: DAQ-related budget request showing number of computers and costs in k$.

The total proposed budget request from US Nuclear Physics funds for CMS DAQ-related activities is estimated to be 2 M$. This contribution from the US heavy ion groups will constitute about 40% of the CMS computing power for the High Level Trigger and Event Filter of CMS, about 10% of the overall cost of the full DAQ system, and about 0.4% of the total cost of the CMS experiment. It should be noted that the CMS budget for future years includes funding for maintenance and gradual replacement of the full HLT, including the portion initially purchased using US Nuclear Physics funding.

3 Operations Budget

The CMS collaboration, following the long established CERN practice, maintains a fund to which all collaborating institutions or countries contribute. This fund is used to provide resources for detector operations at CERN. The size of the contribution to this "M&O Category-A" fund is calculated annually, each September, and is proportional to the number of active PhD physicists in the group. It is important to note that a person needs to be a collaboration member in full standing for at least one year in order to appear as an author on CMS physics publications. Table 2 gives a present estimate of the number of the US physicists expected to work in the CMS heavy ion program. The contributions for the physicists working on the heavy ion program will need to be covered by heavy ion physics funding agencies. It is expected that the contribution per physicist in FY2007 will be about 7.7 k$ without overhead, increasing gradually to about 10.5 k$ for FY2010 and beyond. "M&O Category-B" funds cover the responsibility of groups that have built detectors to maintain their equipment. Currently, the only item in this category for the US CMS heavy ion group is the Zero Degree Calorimeter which will need yearly maintenance with a likely replacement of the tungsten and fibers in FY2010 or FY2011, after a few years of operation. Note that upgrades and maintenance of online computing will be covered using "Category A" funds collected from the entire CMS collaboration.

The US CMS heavy ion group will provide a pair of Zero Degree Calorimeters (ZDC) which will serve a vital role in the heavy ion program at CMS. The ZDCs are mounted behind the first accelerator magnet at the end of the straight section surrounding the interaction point. Located 140 m from the vertex and installed into a space 1 m long and 9.6 cm wide between the two beam pipes, they cover $|\eta| \geq 8.5$ for neutral particles. The calorimeter design is similar to that for the RHIC ZDCs but with the addition of a segmented electromagnetic section which allows the detection of photons as well as neutrons. The Kansas, Iowa, and UIC groups are primarily responsible for the two ZDCs. The radiation load on the phototubes is comparable to that of the tubes for the HF calorimeter. However, it is possible that a beam accident could damage phototubes of one or the other ZDC. For this reason, funds for ZDC maintenance are included in the operating costs discussed below.

The large distance to CERN will result in some increase in travel costs which will vary from group to group. To provide the essential direct contact with experts at CERN, we expect that typically 1-2 physicists from each institution as well as the graduate students will stay at CERN for extended periods of time, totaling roughly 14 FTEs on site. Adjustments for cost of living and currency exchange protection are expected to increase costs about $5-15$ k$ annually per person (lower for students than for PhDs) stationed full time at CERN compared to what was typical at Brookhaven. For the entire US CMS heavy ion group of about 30 PhDs with an additional 15-20 graduate students, we expect that the overall cost of travel will be about 250 k$ more than the cost of the same groups doing research at RHIC.

The total operating budget without overhead for the US CMS heavy ion program, including collaboration fees and the expected increase in travel costs is summarized in Table 2.

Year	PhDs	Students	M&O Cat A	M&O Cat B	Travel	Total
	Total number		Costs in k$			
FY2007	20	10	154	5	200	359
FY2008	30	18	276	5	250	531
FY2009	30	18	306	5	250	561
FY2010	30	18	321	5	250	576
FY2011	30	18	321	30	250	576

Table 2: Increase in operations costs without overhead for the US CMS heavy ion program in k$.

The operations of most of the groups involved in this proposal are currently funded by the DoE Office of Nuclear Physics. The group at the University of California at Davis is funded by the National Science Foundation. The Los Alamos National Laboratory group is currently funded by the lab's LDRD program and is seeking DoE Nuclear Physics funding to continue their involvement afterwards. The University of Iowa group is also seeking DoE Nuclear Physics funding to cover its CMS heavy ion research program.

4 Computing for Heavy Ions

At the top of the hierarchy of the CMS computing model is the Tier-0, located at CERN. It receives the raw data from the detector during running and is responsible for archiving the primary copy of the data to tape. The first pass reconstruction of the events is performed in real time at the large computer farm of the Tier-0. In parallel, the raw data is distributed to the next layer, the small number of Tier-1 centers located in different countries. These centers are responsible for storing a second copy of the raw data. Tier-1 centers also supply processing resources for further reconstruction, calibration, and other data intensive analysis tasks. Batch processing of data analysis as well as Monte Carlo simulations are performed at the larger number of Tier-2 centers which each provide a significant amount of CPU power and disk storage. Finally, the Tier-3 centers provide interactive access to local groups and additional modest computing capacity for the experiment. Users at the Tier-3 centers will rely on Tier-2 for access to the data. Again, it is important to note that the CMS planning makes no provision for analysis resources allocated to heavy ion physics.

The heavy ion data and analysis are sufficiently different that adjustments must be made. As such, it makes sense to place all of the heavy ion data in one or more Tier-2 or comparable centers. The goal is to have a system which is sufficient to process and analyze the collection of data that will accumulate over the consecutive data taking periods. Based on the planned data-taking capacity, data volumes are expected to be ∼250 TByte of raw data, which probably increases to ∼430 TBytes of reconstructed data for each full luminosity heavy ion run. Based on a quote from the Fermilab Tier-1 facility, the cost of tape storage of the heavy ion data, as well as Monte Carlo simulation results, is expected to be 75 k$ annually starting in FY2008.

It is expected that significant computing resources beyond those allocated in the current CMS plan will be needed to carry out off-line data reconstruction and analysis for heavy ion physics. The US groups would like to store all raw heavy ion data locally in order to facilitate full reconstruction and analysis by the faculty and students. Based on experience with the RHIC experiments, it is desirable to have sufficient local storage for the processed data from at least two years of running, in particular to allow efficient analysis of rare events. In order to keep the analysis commensurate with the running schedule, it is necessary to have sufficient processing capacity to fully reconstruct a significant fraction of the data from one run within the one year expected between runs. The cost of providing these critically necessary resources is estimated at 325 k$ annually. Some of the details are given in Table 3. Note that in FY2007, only simulation results will need to be stored, and therefore tape costs are smaller. As a consequence, more money is dedicated to buying CPUs. In each of the following years, a roughly equal number of CPUs will be purchased, while machines that are older than 3 years will be replaced. The total heavy ion computing cost is 400 k$ per year, starting in FY2007 and continuing for the duration of LHC running. This sum includes a small charge for tape storage at the Tier-1 center at Fermilab in FY2007 increasing to 75 k$ annually thereafter.

This budget includes only the tape storage and computing costs and does not cover infrastructure such as power and cooling. Those would need to be provided by the host institution(s) or other funding sources. Possibilities currently being actively pursued include sharing infrastructure with CMS high energy Tier-2 facilities under construction at MIT and proposed at Iowa and Vanderbilt, as well as using matching funds from the State of Illinois towards the construction of a CMS heavy ion computer facility at UIC. The participating groups will continue to work both individually and collectively to explore the availability of additional resources for off-line computing. Additional capacity, especially in the area of CPU power, would allow more flexibility in the scheduling of analyses, more opportunity to explore a broader variety of physics topics, and an improved ability to respond expeditiously to new discoveries.

Fiscal Year	Tape [TB]	Batch Nodes	CPU [SI2k.s.10^{12}]	CPU Total [SI2k.s.10^{12}]	Disk [TB]	Total Disk [TB]
2007*	9	105	20	20	157	157
2008	430	83	23	43	186	343
2009	430	83	33	75	279	623
2010	430	83	47	103	419	884
2011	430	83	69	149	628	1326

Table 3: Computing resource purchase schedule. $*$ Note that the FY2007 entry also includes computers purchased in FY2006. See text for details.

5 Personnel Resources

We estimate that by 2007-2008 the number of US heavy ion participants will total roughly 20 faculty and senior research personnel, 10 postdocs, 20 graduate students, and 10 undergraduate students. The non-US heavy ion contingent in CMS will number about 50 people. The current group members are listed in Table 4.

Institution	Name	Position
U. California at Davis	M. Calderón de la Barca Sánchez	Faculty
	D. Cebra	Faculty
	R. Vogt	Adjunct Faculty*
U. Illinois at Chicago	O. Barannikova	Faculty
	R.R. Betts	Faculty
	E. García	Research Faculty
	D. Hofman	Faculty
	R. Hollis	Post-doctoral
	A. Iordanova	Post-doctoral
U. Iowa	N. George	Adjunct Faculty
	F.D. Ingram	Adjunct Faculty
	E. Norbeck	Faculty
	Y. Onel	Faculty
U. Kansas	O. Grachov	Post-doctoral
	M. Murray	Faculty
	S.J. Sanders	Faculty
Los Alamos Nat. Lab.	P. Constantin	Post-doctoral
	G.J. Kunde	Research Staff
	C. Mironov	Post-doctoral
U. Maryland	A.C. Mignerey	Faculty
	M.B. Tonjes	Post-doctoral
Mass. Inst. of Tech.	M. Ballintijn	Research Scientist
	W. Busza	Faculty
	C. Loizides	Post-doctoral
	C. Roland	Research Scientist
	G. Roland	Faculty
	G.S.F Stephans	Senior Research Scientist
	B. Wyslouch	Faculty
U. Minnesota	J. Kapusta	Faculty*
Rice U.	B. E.Bonner	Faculty
	G. Eppley	Research Staff
	W. Llope	Research Faculty
	G. Mutchler	Faculty
	P. Yepes	Research Faculty
Vanderbilt U.	S.V. Greene	Faculty
	C.F. Maguire	Faculty
	J. Velkovska	Faculty

Table 4: Current institutions and non-student members of the US component of the CMS heavy ion group. Names of theorists participating in the program are marked with $*$.