TWEPP-08 Executive Summary

J. Christiansen [a], P. Farthouat [a], F. Formenti [a], A. Gonidec [a], G. Hall [b], K. Kloukinas [a], J. Oliver [c], E. Petrollo [d], C. De La Taille [e], F. Vasey [a], T. Wijnand [a], K. Wyllie [a], R. Yarema [f]

[a] CERN, 1211 Geneva 23, Switzerland
[b] Imperial College, London, United Kingdom
[c] Harvard University, Cambridge, USA
[d] INFN, Rome, Italy
[e] Laboratoire de l'Accélérateur Linéaire, Orsay, France
[f] Fermilab, Batavia, USA

I. Some Statistics

The Topical Workshop on Electronics for Particle Physics (TWEPP-08) took place in Naxos, Greece, from 15 to 19 September 2008. Thirteen invited and 126 contributed papers (61 oral and 65 poster) were presented in 7 plenary and 11 parallel sessions to an audience of approximately 160 participants. Twenty-one of these participants came from the United States and three from Japan, while the rest originated from Europe. Of all presented papers, 43% referred to the LHC project, 25% to the SLHC upgrade, and 32% to ILC and other experiments.

II. Session Summaries

Some of the main conclusions from the sessions dedicated to **ASICs**, **Optoelectronics, Power, Installation and Commissioning**, **LHC Upgrades, Trigger** and **Posters** are summarized in Sections A to G below. Owing to space constraints, many contributions have to be omitted from this summary, but the interested reader can find them all at [1]. Invited presentations during the opening session of the workshop surveyed the High Energy Physics activities in Greece, highlighting a healthy programme covering accelerator-based physics at CERN, DESY, and FERMILAB, as well as other experiments such as Nestor and KM3NeT. Two other HEP projects were presented: the European XFEL Photon Source and the Large Synoptic Survey Telescope (LSST). Many additional invited talks introduced the sessions reviewed below.

A. ASICs

The two ASIC sessions hosted 12 talks, 9 of which were dedicated to pixel readout chips, showing the important R&D effort needed to meet the challenges of the next generation of tracking detectors and calorimeters. These chips are intended for the SLHC, the ILC or the Gigatracker at NA62. For all these, the challenges of denser integration, minimizing power dissipation and tolerating increased radiation levels are being tackled through new technologies or innovative design. The three remaining talks addressed Si-Photo-Multiplier readout, high speed serializers, and the radiation hardness of SiGe technology.

Electronics designers have to meet new challenges in order to provide larger systems at a lower cost per channel. Modern microelectronics technology makes it possible to implement in silicon more functions than ever before (analog and digital blocks, often combined together with several programmable features). Several technologies are available to the designers. They evolve so rapidly that it is often a challenge to get the best compromise between the constraints of project planning and the choice of the latest technology. Design engineers have to select their preferred technology on the basis of the performance required at the system level, the available resources, the requested delivery time, and the long term availability of the technology. The range of available library functions best suited to the application must also be taken into account.

An important requirement to any chip operating in the front-end of SLHC systems is radiation hardness. Modern fine-feature technologies appear to offer this hardness almost without any special preventative measure; however, engineers continuously looking for best performance at the highest radiation levels must constantly confirm this fact.

Besides coping with the increase in electronic channel density and the need for radiation hardness, achieving higher speed with reduced power dissipation becomes the next design challenge. Modern microelectronic technologies can contribute to solving this problem, thanks to their lower power supply voltage. Generating and distributing this power was the subject of a dedicated session reviewed in section C below.

Testability is an important issue at the system level, but in the ASIC domain an almost costless solution is possible. With the capability to integrate more functions per surface area, designers have the possibility to implement inside the ASIC several complementary functions to check its correct behaviour. This does not substantially increase the cost of the final device, while considerably

decreasing the costs related to debugging and future on-board verification and maintenance.

The Microelectronic User Group (MUG) meeting which followed the ASICs session of the workshop consisted of a presentation of the design tools and of the foundry access services provided by CERN to the HEP electronics design community, followed by an open discussion among the participants. The presentation gave an overview of the technologies that are currently supported at CERN and focused on the CMOS 130nm technology that presently constitutes the mainstream for the designers in the HEP community. The digital design flow and the technical challenges associated with designs in modern deep submicron technologies were addressed. Complex physical design rules and manufacturability design constraints, signal integrity issues caused by signal cross-couplings and tough final timing closure for digital designs, requirements for systematic multiple corner simulations, requirements for rigorous analysis of voltage drops on power distribution lines and of electro-migration effects, can only be tackled effectively by employing modern ASIC design tools. Industry offers sophisticated CAE tools to master the difficulties of modern technologies and the key element for delivering successful designs is now the definition and implementation of a specific work flow with those tools formalizing a design methodology. Adoption of common design methodologies and training are two key elements assisting the work on large distributed designs among multiple institutes, enhancing team productivity and increasing silicon accuracy. A fully automated digital design methodology, packaged in a digital design kit, is currently being employed in the CMOS 130nm technology and has been used to implement a series of digital designs. The digital design kit has been distributed to seven collaborating institutes and five training courses were organized at CERN, in the last two years. This design kit is expected to be replaced by a new one that will cover a broader spectrum of functionalities incorporating a design methodology for Analog and Mixed Signal designs, the most typical case of ASICs in the HEP community. The design kit will integrate the foundry Physical Design Kit with digital standard cell and IO pad libraries and a configuration management tool will automate the setup of the design environment.

The Multiple Project Wafer (MPW) services organized at CERN were also presented. MPW runs help in keeping low the cost of prototype fabrication and small-volume production by enabling multiple participants to share production overhead costs. CERN has also developed a working relationship with the MPW provider MOSIS as an alternative means to access silicon for prototyping when demand is not high enough to justify a CERN organized MPW. With the current industry prices, a CERN organized MPW becomes more cost efficient than a MOSIS run at around 90 mm^2 of silicon area.

During the open discussion session, technical issues with the CMOS 130nm design kit were addressed. The issue of building and maintaining a common IP block repository was raised. The repository can be maintained at CERN and could host IP blocks created by the designers of our community. There was a general consensus that this facility is of great importance. To realize this undertaking there must be a consensus among the designers that some effort will have to be invested from their part. The circuits delivered to the repository have to be packaged as IP blocks to facilitate their integration in the design kit, and must be accompanied by a minimal set of specifications and by a simulation test bench when appropriate. Availability of the designers to provide some level of technical support for their circuit is also important for the success of this endeavour.

The organization of more regular MUG meetings was also suggested, to provide effective means of exchanging information in the community and allow designers to become acquainted with ongoing developments.

B. Optoelectronics

The optoelectronics session of the workshop consisted of 5 oral presentations and 7 posters. In addition, 7 short presentations and discussions took place during the opto-working group meeting which followed the session.

Approximately one third of the contributions were dedicated to existing LHC systems, their commissioning and analysis. As already pointed out in 2007, the optical systems installed at LHC are all operational with a dead channel count well below the 1% level. System-level test and setup procedures as well as tools and results were presented by both ATLAS and CMS, confirming that the good system performance observed had been obtained at the cost of a significant effort during installation and commissioning. It was agreed during the opto working group meeting that these lessons will be appended in 2009 to the existing and published “lessons-learned” report.

The remaining two thirds of the contributions were dedicated to future optical systems. Total dose and total fluence tests of PIN diodes, VCSELs and fibres (both single and multi-mode) were reported. Candidate high speed components were identified which seem to resist SLHC radiation levels. These are all very encouraging results, which will be confirmed and reproduced on a larger scale within the opto-working group. Single Event Upsets generated in PIN diodes were also measured at gigabit per second speeds and analysed statistically. Interestingly, not only zero-to-one bit flips, but also one-to-zero SEU-transitions were observed, as well as multi-bit bursts. In cases where unshielded trans-impedance amplifiers followed the pin diode, very long bursts of errors were also reported, indicating

that not only the photodiode, but also the succeeding analog electronics will need to be appropriately designed and qualified to cope with high particle fluxes. In the field of components and systems characterization, test setups operating up to 10 Gbps were presented. Results were shown, confirming that the know-how and equipment is now becoming available in our community to develop high speed links. A figure of merit was described which allowed a quantitative comparison of transmitters and receivers, paving the way to a broad survey of market components. Finally, a status report of the joint versatile link project was presented, highlighting the progress made during the first six months of work.

C. Power

Power supplies and power distribution are key issues for the next generation of experiments. A lot of work is invested in the subject as shown by the number of presentations and posters in the session. There were 7 oral presentations in the dedicated power session, plus 6 posters and another 5 presentations during the power working group session. The vast majority of them were related to future systems.

The main issue for future experiments is to find an efficient way of bringing power inside the experiment volume where the space is highly limited and where the environmental conditions are extremely harsh (magnetic field and radiation). Some examples of power levels needed for SLHC upgraded trackers were given. They show that while the power per electronics channel is expected to decrease (due to the use of smaller feature size technologies) the total amount of current to be fed to the detector volume is increasing (higher number of channels and lower operating voltage). As more current means thicker cables, solutions must be found to reduce this current and maintain the cable envelope.

Two main routes are being pursued to reduce power supply current: one based on DC-to-DC converters and one based on a serial powering scheme.

The principle of DC-to-DC converters is to transform a high-voltage low-current input to a low-voltage high-current output, the input power being equal to the output power divided by the yield of the converter. This technology is heavily used in commercial equipment but needs some adaptation to our needs. First of all, in order to be effective this solution requires the converters to be positioned inside the detector volume and operate in the presence of a strong magnetic field. Most commercial converters rely on magnetic elements for the energy conversion and these elements will not work at the field levels present in the detectors. There are two ways of attacking this problem: one is to use air core inductors and the other is to use capacitors instead of inductors for the energy transfer. Both options are being looked at and both could be used efficiently. For instance a two-stage conversion scheme was presented in which an inductor based DC-to-DC convertor provides a relatively high current and medium voltage (a few amperes at 3 to 5 V) and then a capacitor based DC-to-DC convertor would be embedded in the front-end chips to deliver the low voltages used by the analog and digital circuits. The overall power dissipation in the system could be highly optimized with such a scheme.

Two other problems need to be studied and solved for the DC-to-DC conversion scheme to be successfully demonstrated. One is the potential noise increase which could come from the switching currents in the air core inductors and the other is the need for operating power devices in a high radiation environment. Several noise studies have been pursued using detector elements from ATLAS or CMS powered with either commercial or custom convertors. This work is very promising. It has led to the definition of a measurement method and reference setup that everybody can use. The second problem is the radiation hardness, where the situation is still unclear. Commercial devices were tested but none of them sustained the required radiation level for the trackers of SLHC. A technology to be used for developing a full custom ASIC had been identified but recent radiation tests have also shown some unacceptable limitations.

For the coming year, work will concentrate on finishing the design of a buck DC-to-DC converter ASIC, on testing several technologies against radiation, and on pursuing the study of commercial devices and of noise effects.

The serial powering scheme is also a DC-to-DC converter but working in a very simple way: N elements needing a current I under a voltage V are put in series and powered with a voltage N*V and a current I. Several measurements have been made, showing that this scheme works and does not introduce extra noise. It also has the advantage that the circuitry required does not need a special technology and hence the radiation hardness is less of an issue than for the DC-to-DC converters. However, at the system levels it presents a number of challenges. Firstly, in a distribution chain the devices are all at different potentials leading to constraints on the connection of the readout electronics to the detectors and the way the readout is implemented (AC coupling becomes mandatory, for instance). Secondly, this scheme can only remain power-efficient if all the devices to be powered consume the same amount of current. Thirdly, the control of the system (switching ON and OFF for instance) has to be studied with great care (when a device is OFF, the current is still flowing through it). Finally, the separation of power supplies on a device (for instance if one wants to split the analog and the

digital supplies of a device) is less efficient power-wise as this can only be done with local linear voltage regulators.

Several successful tests have been done with up to 30 devices serially powered and some custom circuitry is expected to become available soon: a new silicon strip readout chip (ABCn) developed for ATLAS contains the necessary elements for implementing the serial powering scheme and an ASIC designed at Fermilab (SPI, or Serial Powering Interface) provides the shunt regulator and the control logic (including protection circuitry) necessary for serially powering a device.

For the time being, it is too early to choose between these two approaches and the R&D work has to be pursued in parallel. It is important, however, to put in place all the methods allowing coherent testing and an easy comparison of performance between the different solutions. A good fraction of that is already in place (e.g. the noise characterization) and it was agreed during the power working group session to design the readout hybrids for the ATLAS R&D work on strip detectors with the capability of including different power devices.

D. Installation and Commissioning

The successful startup of the LHC on 10 September, just before the beginning of the workshop, was a major milestone for the accelerator and all experiments. The establishment of a circulating beam in the presence of media from all over the world was the result of carefully carried out hardware and sector commissioning and of various successful synchronization tests with beam in the preceding months. Both machine and experiments worked extremely hard to meet the tight schedule as exemplified by several presentations:

- The accelerator cryogenic control system, for example, is partially located in the tunnel and is therefore designed to be radiation tolerant. To keep the magnets in nominal conditions, the distributed cryogenic instrumentation connects to approximately 18000 sensors and actuators. To commission the tunnel electronic crates, a Mobile Test Bench (MTB) based on a PXI platform was developed. The MTBs are equipped with a Software Configuration Management (SCM) tool that provides a centrally managed storage area.

- For the ATLAS pixel detector, dedicated calibration techniques were developed for each detector assembly stage, matching the demands of the real detector services and readout system. Nearly the entire pixel detector was successfully commissioned in situ which allowed combined cosmic runs with the rest of the tracker (SCT and TRT).

- From 2006 to 2008, the full detector system for the ALICE silicon pixel detector (SPD) was tested in a dedicated area at CERN. At the same time, as much as possible of the read out electronics were commissioned. Together with the use of automatic control procedures, this enabled the successful recording of the first particles generated by injection tests in the LHC in June 2008. The SPD has been operated almost without interruption since that time.

- The CMS high resolution electromagnetic calorimeter (ECAL) on- and off-detector electronics were commissioned within CMS in situ. An important milestone was the startup of the monitoring system that uses the beam abort gap to calibrate the system. During the synchronization tests of the LHC ring, 98% of the crystals lit up when a beam was dumped on a collimator in the tunnel upstream. ECAL is now operating as expected with occasionally some problematic channels.

- The procedure for the hardware installation and commissioning of the ATLAS LAr calorimeter system was to perform a test as soon as possible in standalone mode. If successful, the elements were integrated with the previously commissioned detector and continuously exercised with pedestal runs, calibration runs and cosmic runs to verify its behaviour and stability. The detector now has 100% of the HV channels operational with a very stable readout system (calibration constants stable to better than 0.1% over a few months).

In summary, the installation and commissioning session was the occasion to review the lessons learned so far. After completion of the commissioning of the individual systems, global tests with cosmic rays were successfully performed and concluded with the observation of the first tracks from the beam halo generated by the first circulating beams in the LHC on 10 September. The cosmic runs provided a useful test of the software, of the readout chain and of the detector performance in the global environment and in the presence of cross-talk and noise. They also provided a test of the detector alignment. Various difficulties encountered during the individual and global commissioning were reported and lessons for the future were pointed out. Despite difficulties which, in some cases, caused important delays (cooling problems, noise problems during large scale tests) most of the detector systems have been successfully tested to perform to the expected performance levels. The LHC experiments are now looking forward to the first collisions.

E. LHC Upgrades

The sub-title of this year's topical session was "needs and reality", quite vividly illustrated in descriptions of requirements for ATLAS and CMS operation at 10^{35} $cm^{-2}s^{-1}$ luminosity. The LHC machine will be upgraded in two stages, and may reach a luminosity ~3×10^{34} $cm^{-2}s^{-1}$ about five years after the LHC switches on, once collimator and final focus quadrupole improvements have been made. This will already require greater performance from the detectors, surpassing that of the current designs.

Beyond the first machine upgrade phase, the picture is less clear but ideas are being explored to find the best way to achieve an order of magnitude increase in luminosity as early as 2018 following a lengthy shutdown during which detector installation could take place. The physics goals depend partly on discoveries in the coming years but it is expected that demands for a ten-fold luminosity increase will be justified in a few years. In this new environment, detector performance will need to be maintained despite pile-up.

Some machine improvements involve higher beam currents, higher reliability or shorter fill time but others require new machine elements or ideas such as magnets inside the experiments for "early separation" of proton beams. This would have a major impact on detector design and likely backgrounds. Others, such as crab cavities, might be easier to adapt to but remain unproven. In most scenarios, luminosity levelling to avoid dramatic differences in event rates is considered desirable for best operation, including easing the design of trigger and readout systems.

Detector plans for Phase I are reasonably clear: incremental improvement to DAQ systems and a 6–8 month shutdown to replace the radiation damaged innermost tracking detectors. In CMS the pixel detector can be replaced quickly, while ATLAS requires at least one year to replace the innermost B-layer. Therefore, plans to insert a new B-layer inside the current detector, along with a smaller diameter beam pipe, are being considered. In the longer term, improved triggers at Level-1 with new contributions from tracking detectors, or topological triggers combining different sub-detectors, are expected.

A vital aspect of future tracking detectors is cooling, especially given expected power challenges. It is envisaged that more heat must be removed than at present, however carefully future electronics are designed. Requirements for silicon detectors are clear: many distributed heat sources in large volumes, low temperature gradients between them, and constant cooling avoiding thermal runaway of irradiated silicon. Cooling pipes must have low mass and a low structural impact, and the fluid should be radiation resistant. The use of evaporative carbon dioxide cooling, as used in a pioneering development for the LHCb Vertex detector, has attracted a lot of attention as a promising candidate for future cooling.

Minimizing front-end electronic power consumption is vital to all tracker upgrades. New IC technologies offer significant power savings per channel but the number of channels is expected to increase significantly and digital, rather than analog, power consumption may become dominant so techniques to minimize it must be exploited. Power savings can also be achieved by adopting low current differential signalling schemes to replace present standards. Prototypes developed for a possible upgrade of the CMS pixel detector have demonstrated less than 10 pJ per bit per link at 160 Mbit/s over a 2 m long low mass unshielded twisted pair cable.

The ATLAS Pixel Detector upgrade envisages a larger system with more layers. New front-end chips must be developed, with a new module, readout architecture and powering scheme, aiming to reduce the material budget. New optical links with higher transmission speed are required, and a new detector control system is also needed.

The architecture of the readout electronics for the upgraded ATLAS tracker is a change from the current design. The detector is organized in staves with a hierarchical readout with fewer but higher speed links. Some components are not required in very high quantities which may encourage common solutions. Preliminary studies of a front-end preamplifier-shaper-discriminator in 130nm CMOS show a power dissipation below 200 μW per channel while its distribution requires special efforts; savings of factors 5–10 in total current compared to the present system are thought possible.

Progress towards a new CMS microstrip tracker readout for SLHC includes evaluating advantages and disadvantages of different architectures, including trade-offs between power, FE chip and system complexity, system robustness, and performance. A three year development programme is about to start, to deliver a full chip prototype in the second year. Several system level decisions are still open, such as sensor choices, powering scheme and choice between analog and binary, sparsified or not. There are significant advantages in using a binary, non-sparsified design for short strip readout, which results in a simpler chip and system. However, triggering remains the most challenging aspect of the CMS tracker for SLHC, and ideas are still developing which require more simulation results.

Upgrade plans are not restricted solely to the trackers. A possible upgrade of the ATLAS Monitored Drift Tube Detector was presented, allowing it to cope with SLHC background rates. There are several options: to use small tubes, or tubes with field shaping, or other gaseous detectors, such as MicroMegas. For the hottest regions of the detector, 150–180 MDTs will have to be rebuilt covering 600–700 m^2. Either an increased global readout bandwidth by factor 5–10 must be used, or selective readout based on trigger chamber information. For less hot regions, FE electronics should also be redone, ideally using radiation tolerant FPGAs. Everywhere, ageing behaviour must be evaluated and improvements in the radiation tolerance of power supplies achieved.

Peter Sharp described in his invited talk some lessons learnt during the development, production and commissioning of the LHC detectors and their electronics. He summarized his experience and recommendations to the HEP community on microelectronics, optoelectronics, power supplies

and off-detector electronics, all of which raised many important issues for future system development. It is widely acknowledged today that effective use of microelectronics is vital for high performance particle physics detectors. In the early days of the LHC, microelectronics was still emerging and significant investments in skilled manpower and tools were needed to take advantage of modern CMOS technologies. This was accomplished taking advantage of support for microelectronics teaching in universities via EU-funded projects, allowing affordable tools and access to technology.

Appropriately skilled personnel are vital, becoming even more so as technologies advance, and rely on increasingly complex CAE tools. To work in the very hostile radiation environments of the LHC with minimal power consumption requires additional expertise to be developed and maintained. Partnerships with foundries have been found important. For development of systems for SLHC, design groups with sufficient critical mass to master new technologies and tools will be critical. The increased complexity and cost of ASIC developments will surely imply that effort must be concentrated on a limited number of different front-end chips to be produced in large quantities.

Optical link technology is also now vital and the development of radiation hard links has required a large qualification effort, where many details contribute to ensure highly reliable systems. Links with very different characteristics have turned out to have very similar production and development costs, in contrast to some expectations.

Power supplies for LHC front-end systems provide an example where it would have been advantageous to include them earlier in system designs. Supplying clean power and implementing grounding and shielding is non-trivial. Constructing systems to supply power over significant distances and tolerating radiation and magnetic field in experimental caverns required non-standard industrial products.

Off detector, many LHC systems profited from the performance and flexibility of FPGAs but in some cases choices were fixed before FPGAs were widespread. Clearly, technology decisions should not be made too early to profit from commercial electronic progress. Extensive use of FPGAs has delivered benefits but also introduced problems of managing large amounts of complex firmware that must be maintained, supported, and probably upgraded over time.

In future, the community must find (more) efficient ways to produce systems with hundreds of thousands of chips, tens of thousands of modules and optical links, thousands of power supplies and large, complicated readout and DAQ modules. Engineering and quality assurance of such large systems will require a disciplined and professional approach and may require changes to the way HEP collaborations operate. The expertise to build and run such systems must be distributed so that large projects do not become over-dependent on a few key individuals. It is vital to maintain a critical mass of experienced engineers, with appropriate facilities, to conserve experience from current systems.

The discussion session which followed the presentations had a lively debate on the question of appropriate time schedules. It is clearly acknowledged that such large and complex systems require significant time for their development, production, qualification, integration and commissioning. For Phase 1 upgrades there is practically no time to do basic R&D. In important cases it is necessary to start development of Phase 2 SLHC upgrades systems soon to be capable of having working systems available in ~10 years. This may leave very little time for generic R&D programmes based on new technologies that potentially offer major advantages to new detectors.

Spending too much time on generic R&D generates a risk of not having systems ready in time. To take advantage of new technologies it will be important to find ways of shortening the time needed for production, integration and commissioning. This is clearly a major challenge for large collaborations and expected to be a major subject for discussion over the next few years.

F. Trigger

The Trigger session presentations were largely focused on the commissioning and first operation phase of the trigger systems used as part of the regular data taking: recording cosmic muons for long periods and comparing rates with Monte Carlo expectations and measured flux in the cavern. Whilst waiting for first collisions, the focus has been mainly on the development of timing and energy calibration procedures.

During the commissioning phase, misconnections of cables and bad modules were fixed and intensive tests of signal integrity and stability were performed. The challenging requirements at the board and system level could then be verified. The accumulated experience will allow for a better control of the systems in terms of stability and trigger rates, and to fulfil their main role of providing a reliable trigger decision to the experiments.

The work of the present and past year was also largely dedicated to the development of software tools essential to system operation and monitoring. On the hardware front, the extensive use of FPGAs made it possible to adapt board behaviour to system performance by just changing and upgrading the firmware.

Most recently, the systems have been successfully used with LHC related particles during

the very first tests, in injection and dump mode and with single beam radio-frequency capture, collecting events reconstructed in the detectors.

Studies on SLHC trigger systems are ongoing. Much work still needs to be done in the definition of trigger requirements and in identifying the appropriate technology for system implementation.

G. Posters

The poster session contained some sixty-five contributions covering a broad range of topics including ASICs, Trigger, Power Conversion, Optoelectronics and Experiment Commissioning. Just over half of these covered developments for LHC experiments, with the others representing a broad range of interesting applications. Overall, the depth and quality of the poster presentations was impressive. The session enabled in-depth one-on-one discussions in a way which is generally not possible in oral sessions. It was very well attended and generated lively discussion between the attendees and poster presenters. Just over half of the posters were presented outdoors where the balmy Aegean breezes enhanced the experience.

III. CONCLUSION

The TWEPP-08 workshop was characterized by a very open and positive atmosphere. It provided the right framework for discussing old and new projects, reviewing developments, encouraging robust engineering and quality assurance practices, and discussing access to technologies and tools. The oral and poster sessions, together with the parallel working group meetings proved to be the right forums to ensure that critical mass is maintained in our community and that lessons learned are passed on from generation to generation.

Finally, the local organizing committee must be praised for turning this workshop into such a pleasant event [2] in a truly exceptional location.

IV. LINKS

[1] http://indico.cern.ch/event/twepp08
[2] http://www.ntua.gr/twepp08/