# 3rd ATTRACT TWD Symposium in Detection and Imaging (Tripolis)

Wednesday, 31 May 2017 - Thursday, 1 June 2017

Other Institutes

# Book of Abstracts

# Contents

**82**

# Node-X: A networked architecture for energy efficient high performance computing and data acquisition

**Authors:** Evangelos Angelakos[1]; Spiros Poulis[2]; Grigoris Dimitroulakos[2]; Konstantinos Masselos[2]

[1] *Nanotronix Inc - University of Peloponnese*

[2] *University of Peloponnese*

**Corresponding Authors:** kmas@uop.gr, pcst15001@uop.gr, pcst14009@uop.gr, dhmhgre@uop.gr

Introduction
Energy dissipation currently represents the major bottleneck towards exascale computing. High end instruction set processors including many/multi – core processors and graphics processing units (GPUs) adopt an instruction stream based/temporal model of computation which is energy inefficient. Field Programmable Gate Arrays (FPGAs) and Application Specific Integrated Circuits (ASICs) allow implementation of a data driven/spatial model of computation which is more energy efficient. Despite interconnect overheads FPGAs are about 4 times more energy efficient than state-of-the-art many/multi-core CPUs and GPUs.
Different options exist for interfacing accelerator components (such as FPGAs and ASICs) to host CPUs. These include tightly coupled approaches such as integration of CPU and accelerator in a single integrated circuit and CPU motherboards with FPGA sockets. Such approaches are cost inefficient and present serious scalability constraints as far as the number of accelerator components over host CPUs employed is concerned. Expansion buses such as PCIe bus form an attractive choice for loosely coupled accelerators due to their high throughput capacity. However such approaches are expensive and require significant code development efforts both on the accelerator (FPGA) and host CPU side while scalability remains an issue (the number of PCIe slots in a CPU motherboard is limited).
Node-X computing architecture aims at overcoming limitations of state of the art computing architectures in terms of energy efficiency, scalability, cost and programmability.

The Node-X architecture
Node-x computing architecture deploys FPGA/ASIC-based, stand-alone acceleration boards and uses conventional/high-speed Ethernet to interconnect them. Limiting per-node processing elements (FPGAs or ASICs) to one, 'Node-X nodes' build highly granular, Ethernet based, computing fabrics.
Early work on analyzing data-flow patterns in modern data-centers showed that, in certain (common) processing scenarios, moving data through Ethernet (between processing nodes) is not introducing energy consumption or latency penalties given the distributed nature of client-server/API based programming paradigms becoming dominant. Additionally, 4x10G Ethernet interconnected nodes were measured to process and round-trip data in more predictable and efficient ways (compared to locally processed data), without evidently stressing per-unit (node) Bill of Materials (BoM), given the popularity and availability of mature 10G Ethernet silicon.
Node-X nodes are architectured to operate in Single Input Single Output (SISO) mode, foreseeing two possible hardware implementations:
i. Processing nodes: Collect input data from Ethernet – process using FPGAs or ASICs – and release output again on Ethernet and
ii. Data provision nodes: Source/Sink data between Ethernet and Memory or I/O
Abstracting data provisioning from processing nodes allows data residing in volatile/non-volatile memories to be handled in equal footings with data coming from high-speed I/O (ADC/DAC), making node-X an ideal candidate for DAQ applications while alleviating programming from handling data concurrency challenges.
Furthermore, by having node-X nodes built around a single FPGA or ASIC (per node), further gains in lifting programming complexities and in increasing performance predictability are observed, together with critical benefits in assisting fault tolerance scenarios to be deployed, enhancing hardware level granularity, as well as sizing-down per-node manufacturing costs.

**Signal processing, data acquisition**:

**System integration and engineering**:

**Computing**:

**Software and imaging**:

**85**

## Introduction - What is ATTRACT?

**Corresponding Author:** sergio.bertolucci@cern.ch

**75**

## Keynote talk - Bright future for Hybrid Pixel Detectors

**Corresponding Author:** christian.broennimann@dectris.com

**89**

## Presentation 3 - Integrated control and data acquisition systems

**Corresponding Author:** enrico.gamberini@cern.ch

**79**

## Integrated control and data acquisition systems

**Authors:** Giovanna Lehmann Miotto[1]; Enrico Gamberini[1]; Roland Sipos[2]

[1] *CERN*

[2] *Eotvos Lorand University (HU)*

**Corresponding Authors:** giovanna.lehmann@cern.ch, roland.sipos@cern.ch, enrico.gamberini@cern.ch

Traditionally there has always been a rigid separation between control and Data AcQuisition (DAQ) systems for detectors. This comes from the fact that historically very different technologies were applied to these types of systems and that the main system requirements, i.e. high reliability vs. high speed, were considered incompatible. The situation is different today, and this proposal aims at exploring a much tighter integration between control and data acquisition.

The starting point for this study derives from the wish to minimize the number of physical links on the on-detector electronics, to limit cabling effort and dead-material. This is particularly relevant for the inner layers of the experiments at the LHC, but it is an equally important requirement also in other domains. At CERN, a bi-directional protocol (GBT) has been developed for this purpose, and is going to be used in many detector upgrades.

Once the physical layer is shared in close vicinity to the detector, it becomes natural to start thinking on how the downstream part of the DAQ and control systems may be further integrated as well. To this purpose, an architecture as proposed in Abstract 78 allows to route all data to/from the detector to different end-points sharing a common network. DAQ and control may thus just be different peers on the network. Certainly, carefully applied Quality of Service (QoS) methods on the network, in conjuntion with the selection of appropriate communication protocols, are important ingredients to the success of such integration.

Last but not least, higher levels of the software could also be better harmonised, by adopting and sharing similar technologies at the supervisory layers of the control system and the run control of the DAQ.

If successful, the proposed study will lead to the development of a highly integrated DAQ/control system that may be adopted in a broad range of environments, thus reducing the duplication of effort and allowing the experts to focus on the detailed aspects of their own setup and final goals.

**Signal processing, data acquisition**:

data acquisition, control systems, modular design, supervisory software

**System integration and engineering**:

integration of data acquisition and control

**Computing**:

**Software and imaging**:

**86**

## Keynote Talk - Where is the so-called “plenty of room at the bottom”?

**Corresponding Author:** pablo.garcia.tello@cern.ch

**1**

## Keynote Talk - implications of the advances in detector and imaging technology for generating scientific ‘Big Data’

**Corresponding Author:** tony.hey@stfc.ac.uk

**14**

## Summary conclusions of Day 1

**Corresponding Author:** krisch@esrf.fr

**4**

## Presentation 2 - Modular data acquisition systems centred on commercial networks and compute nodes

**Corresponding Author:** giovanna.lehmann@cern.ch

**78**

## Modular data acquisition systems centred on commercial networks and compute nodes

**Authors:** Giovanna Lehmann Miotto[1]; Roland Sipos[2]; Enrico Gamberini[1]

[1] *CERN*

[2] *Eotvos Lorand University (HU)*

**Corresponding Authors:** roland.sipos@cern.ch, giovanna.lehmann@cern.ch, enrico.gamberini@cern.ch

With the increased need for scalable data acquisition systems, we see high potential in modular designs not only for software, but also for the hardware elements. Oftentimes, the on-detector electronics sends and receives data over point-to-point links which are connected to custom processing units. This proposal aims at minimizing the need for custom devices, by introducing an intermediate layer, based on commercial components, that will decouple the point-to-point links from the data consumers/producers. This allows the latter to be implemented by software in commodity servers attached to the network. Such an approach is being actively researched in ATLAS for the future LHC upgrades, but is applicable more generically to many data acquisition and high performance computing architectures.

Making use of FPGAs at the decoupling layer level allows to support multiple point-to-point protocols, depending on the detector needs and its environment (e.g. radiation levels, data volumes, synchronization, reliability, etc.). Integrating the FPGAs into commodity servers (e.g. via PCIe) allows making use of large amounts of cheap memory for buffering data, in combination with substantial computing resources to orchestrate the routing of data to/from the detector to the relevant end-points via high-speed networks.

Our goal is to provide the enabling of a wide variety of users to build data acquisition systems from these building blocks with more streamlined R&D efforts, and to emphasize the importance of scalable, flexible, uniform and upgradable TDAQ systems.

**Signal processing, data acquisition**:

FPGA, high-speed networks, modularity, high performance computing

**System integration and engineering**:

modular design, scalability, use of commercial components.

**Computing**:

**Software and imaging**:

**2**

## Presentation 1 - Node-X: A networked architecture for energy efficient high performance computing and data acquisition

**Corresponding Author:** kmas@uop.gr

**81**

## Intelligent Decision Support System for Environmental Risk Management

**Authors:** Panagiotis Nikolaou[1]; Kostas Peppas[1]

[1] *University of Peloponnese*

**Corresponding Authors:** pananiko2@teiemt.gr, peppas@uop.gr

IDSERM: Intelligent Decision Support System for Environmental Risk Management

Wildfire activity is strongly affected by four factors: fuels, climate weather, ignition agents and people. Where fuel is available, weather is the most important factor in shaping fire regimes in many areas of the world.

In 2009, ISO published the international standard "Risk Management – Principles and Guidelines (ISO 31000:2009)". The ISO risk management standard, which provides generic guidelines for the management of all forms of risk, is being adopted by countries as their national standard for risk management. For example, in Australia and New Zealand, the standard has been published as AS/NZS ISO 31000:2009.
Currently there is no internationally agreed risk management standard covering all types of fire associated with forest and rural land management.
Risk is the "effect of uncertainty on objectives" and this can have positive and negative consequences
Risk Management is the "coordinated activities to direct and control an organization with regard to risk"
Risk Management Process is the "systematic application of management policies, procedures and practices to the tasks of communication, consultation, establishing the context, identifying, analyzing, evaluating, treating, monitoring and reviewing risk"
Risk Management Framework is the "set of components that provide the foundations and organizational arrangements for designing, implementing, monitoring, reviewing and continually improving risk management throughout the organization"
Risks: combining the consequences of a hazard with the likelihood of its occurrence

According to ISO 31010, risks are the combination of the consequences of an event or hazard and the associated likelihood of its occurrence. Consequences are the negative effects of a disaster expressed in terms of human impacts, economic and environmental impacts, and political/social impacts. More detail on the measurement of impacts will be provided separately in the next chapter below.
In situations where the likelihood of occurrence of a hazard of a certain intensity can be quantified we refer to the term probability of occurrence36. When the extent of the impacts is independent of the probability of occurrence of the hazard, which is often the case for purely natural hazards, such as earthquakes or storms, risk can be expressed algebraically as:

Risk = hazard impact * probability of occurrence.

Where the size of the impact influences the likelihood of occurrence, i.e. where the two terms are not independent of each other, the risk cannot be expressed simply as a product of two terms but must be expressed as a functional relationship. Likewise, where the impacts are dependent on preparedness or preventive behavior, e.g. timely evacuation, there are advantages in expressing the impact indicator in a more differentiated manner. In particular in the analysis of natural hazards, impacts are often expressed in terms of vulnerability and exposure. Vulnerability V is defined as the characteristics and circumstances of a community, system or asset that make it susceptible to the damaging effects of a hazard.

Risk=f(p$E$V)

Using the concept of vulnerability makes it more explicit that the impacts of a hazard are also a function of the preventive and preparatory measures that are employed to reduce the risk. For example, for a heat wave hazard it may be the case that behavioral preparedness measures, such as information and advice, can critically reduce the vulnerability of a population to the risk of excess death. Effective prevention and preparedness measures thus decrease the vulnerability and therefore the risk.

Depending on the particular risk analyzed, the measurement of risk can be carried out with a greater number of different variables and factors, depending inter alia on the complexity of the chain of impacts, the number of impact factors considered, and the requisite level of precision. Generally, the complexity of the modelling and the quantification of factors can be increased as long as this also improves certainty.

Motivated by the above facts, in this work, an integrated hardware/software system for environmental decision support has been developed.
The system comprises a weather station and a wireless sensor network based on the development boards Arduino Uno Rev 3 and Raspberry Pi 2. The sensor network employs the 802.15.4 protocol for wireless weather data transmission via XBee S1 transceivers and the IoT publish/subscribe protocol MQTT (MQ Telemetry Transport) for gateway communication. The proposed system estimates the forest fire risk based on statistical analysis of weather data from the environment to which a weather station is mounted. Then, the system computes the Fire Weather Index (Fosberg Index) and the Fire Risk Indicator by employing advanced artificial intelligence techniques, such as fuzzy logic, machine learning and predictive risk models, neural networks and deep learning techniques. The proposed system configuration is useful to engineers and scientists that actively participate in monitoring, detecting and predicting forest fires.

**Signal processing, data acquisition**:

**System integration and engineering**:

**Computing**:

**Software and imaging**:

**88**

## Presentation 5 - Intelligent Decision Support System for Environmental Risk Management

**Corresponding Author:** peppas@uop.gr

**83**

## Crowd-Dx(TM): Using crowd-sourced data analytics to elucidate trends in healthcare, wellness and disease

**Author:** David Sarphie[1]

**Co-author:** Rubina Mian [2]

[1] *Bio Nano Consulting*

[2] *Oxford MediStress Ltd*

**Corresponding Authors:** rubina.mian@oxford-medistress.com, david.sarphie@bio-nano-consulting.com

Introduction:
Developments in Connectivity, Big Data and Structured Data Analytics promise significant advances in many important societal areas, including healthcare. Indeed, a number of predictions about future technological advances relate to "Smart Healthcare", wherein the connected world and Internet of Things (IoT) will impact our understanding of disease and offer material benefits to patients and healthcare providers. A key requirement to make this happen is availability of meaningful and robust data regarding relevant biomarkers that can readily be obtained from a significant portion of the population to enable elucidation of important healthcare trends. Whereas new "wearable" technologies can provide information regarding certain indicators (activity, heart rate, sleep patterns, etc), they are not able (at least yet) to provide data regarding immunological biomarkers. This is a key unmet need that this project will address.

Background to project:
Recent advances in immunology and diagnostics have enabled the development of a sophisticated

home/self-test for a key biomarker of wellness and disease. Based on the function of neutrophils (a type of white blood cell), this rapid blood test has been shown previously to provide important health and wellness data in several applications, including: i.) athlete performance & recovery data for elite sports teams (e.g. Manchester United in the EPL, Maclaren in Formula-1 racing), ii.) data providing a direct indication of stress response for individuals suffering from clinical symptoms of psychological/emotional stress, and, iii.) important new diagnostic/prognostic data for oncologists monitoring disease progression in prostate cancer patients (this from recent clinical trials at UCLH in London). Several clinical advisors to this project have indicated that the clinical biomarker that is measured in this Leukocyte Coping Capacity(TM) (LCC) test is meaningful not only on an individual basis (especially when monitored longitudinally over time) but also potentially when analysed across groups of individuals (population groups) to understand immunological differences between groups which can vary over time. Data from a broad spectrum of the population, when analysed using sophisticated data analytics, could, for example, help healthcare providers spot trends in immunological changes related to health & wellness, stress or infectious disease (for example, in a disease outbreak). Because of its ease-of-use, its potential as a self-test, and its ability to provide key immunological data that impacts large swathes of the population, project partners are confident the LCC test could provide crucial, clinically-meaningful data that will help individuals live better and help healthcare providers spot important trends in the future.

Project description:
In this project the teams will develop the initial data analysis tools and software-based algorithms that will enable LCC data from numerous individuals to be analysed and the data portrayed in a meaningful way. A first-generation smartphone mobile application will be developed. The team will undertake early-stage discussions with healthcare providers to ensure the data analysis results in useful outputs.

**Signal processing, data acquisition**:

**System integration and engineering**:

**Computing**:

**Software and imaging**:

**90**

## Presentation 6 - Crowd-Dx(TM): Using crowd-sourced data analytics to elucidate trends in healthcare, wellness and disease

**Corresponding Author:** david.sarphie@bio-nano-consulting.com

**91**

## Presentation 7 - quAffirm-GPS™: Developing a network of smart sensors & data analytic strategies to address a global environmental issue: excessive arsenic in drinking water

**Corresponding Author:** david.sarphie@bio-nano-consulting.com

**92**

## AquAffirm-GPS™: Developing a network of smart sensors & data

# analytic strategies to address a global environmental issue: excessive arsenic in drinking water

**Author:** David Sarphie[1]

[1] *Bio Nano Consulting*

**Corresponding Author:** david.sarphie@bio-nano-consulting.com

Problem Description:

With one in ten people worldwide lacking access to clean water, the World Economic Forum in 2015 named the water crisis as the #1 global risk based on impact to society "as a measure of devastation". A key contributor to this overarching issue is the problem of chronic exposure to arsenic in drinking water, which affects over 140 million people in 70 countries, leading to a variety of health issues, including skin lesions, cancers, and, in many cases, death. Fundamental to tackling this issue is having access to dynamic relevant data about arsenic levels, which are labile, in individual wells, particularly in highly-affected areas, such as parts of Bangladesh (where the WHO called it the "largest mass poisoning in history"), India, Vietnam, China, Argentina, Chile, Mexico and the USA. Management of this problem is hampered by disadvantages associated with the lab-based chemistry assays or the colorimetric, non-digital point-of-use tests in current use: these are slow, difficult to use, and produce non-digital, single-measure outputs which cannot be easily aggregated into a single regional database. What is needed are rapid, easy-to-use sensors producing digital output that can be connected using rapidly-evolving networks to a centralised database so that modern data analytics can be used to address and manage the problem more effectively.

Background:

Bio Nano Consulting (BNC), together with its academic partners at Imperial College, have developed a unique arsenic sensor – the AquAffirm™ system – that uses technology modified from familiar glucose self-tests manufactured in the billions each year to create a step-change improvement in arsenic testing. By solving customer usability, cost and accuracy needs, which are unmet by current arsenic tests, BNC has developed an innovative platform which enables the commercialisation of simple, rapid, affordable tests for use in addressing this important problem. Crucially, these tests produce digital outputs which can be networked to central databases. A current R&D collaboration with two research organisations in Mexico (funded in part by Innovate-UK, the UK's innovation agency) is delivering a mobile app and the software to enable the test data to be communicated from the bespoke electronic meter the team is developing to a smartphone. This app will include functionality that allows the exact location (within approximately 8 meters, obtained via the smartphone's GPS chip) to be transmitted together with the output of the AquAffirm™ sensor's reading to a central database to enable advanced data analysis.

Project Description:

In the proposed project BNC will develop the initial data analysis tools and software-based algorithms that will enable arsenic data from individual AquAffirm™ tests to be aggregated and analysed and the data portrayed in a meaningful way. Sophisticated data analytics will be developed to enabled the data to be "mined" for meaningful trends to improve the way the project is managed. The team will undertake early-stage discussions with non-governmental agencies (NGOs) such as UNICEF and government agencies such as Bangladesh's Department of Public Health Engineering (DPHE), both of which have been tasked with managing the problem, to ensure the data analysis results in useful outputs which will assist in managing the problem both globally and in the countries most affected.

**Signal processing, data acquisition**:

**System integration and engineering**:

**Computing**:

**Software and imaging**:

Data analytics

**87**

## Presentation 4 - Silicon Photonic Data Transmission for Detector Instrumentation

**Corresponding Author:** marc.schneider@kit.edu

**80**

## Silicon Photonic Data Transmission for Detector Instrumentation

**Authors:** Marc Schneider[1]; Djorn Karnick[1]; Lars Eisenblätter[1]; Yunlong Zhang[1]; Julius Hartmann[1]; Thomas Kühner[1]; Marc Weber[1]

[1] *KIT - Karlsruhe Institute of Technology (DE)*

**Corresponding Authors:** marc.weber@cern.ch, lars.eisenblaetter@kit.edu, thomas.kuehner@kit.edu, julius.hartmann@kit.edu, zhang.yunlong@partner.kit.edu, marc.schneider@kit.edu, djorn.karnick@kit.edu

In future detector systems for particle physics and photon science the ever-increasing number of electronic channels results in a keen demand for data read-out capacity. The data transmission from the detector front-end to the counting room relies more and more on optical links. In state-of-the art solutions, individual optical fibers connect directly modulated laser diodes to a corresponding receiver in the periphery of the detector.

Recently, we proposed an optical data transmission system based on wavelength division multiplexing (WDM) [1, 2], where a single optical fiber transports numerous optical channels, which may provide a generous data rate up to the Tbit/s range. That would increase the data readout capacity significantly while reducing the number of individual fibers connecting read-out chips with the data acquisition units. Laser sources providing the optical carriers would be located off-detector thus not contributing to the energy budget within the detector volume. WDM systems are well-known in telecommunications for long-haul networks but are entirely new to high energy physics experiments. The achievable bandwidth together with low mass and power consumption would revolutionize data transmission, data acquisition, and trigger architectures.

The essential building block of our proposed system is the monolithically integrated multi-wavelength transmitter on a silicon-on-insulator (SOI) platform. An optical demultiplexer separates incident optical carriers in order to forward each of them to an electro-optic modulator to encode electrical information. A multiplexer merges all data-carrying signals to be transported over a single optical fiber. The (de-)multiplexers would be implemented as planar concave gratings (PCG), the electro-optic modulators as Mach-Zehnder interferometers consisting of two identical pn phase-modulators.

For a first conservative demonstrator, we aim for a data rate of 10 Gbit/s per modulator to keep costs and complexity for the driver circuity low. The integrated transmitter would encode data on four WDM-channels, which results in a data rate of 40 Gbit/s per fiber. The per-channel data rate can be increased to 32 Gbit/s by using the full speed of the driver chip and carefully optimized circuitry, adding up to 128 Gbit/s per fiber. Future versions of the optical link could operate 16 to 64 wavelength channels, making a total transmission bandwidth up to the Tbit/s range feasible.

We present our current work on photonic chip design and prototypes, driver design, simulations, and packaging, as well as challenges to enable the system to work at extreme radiation levels.

[1] P. Skwierawski et al., 2016, JINST 11 C01045

[2] D. Karnick et al., 2017, JINST 12 C03078

**Signal processing, data acquisition**:

**System integration and engineering**:

**Computing**:

**Software and imaging**:

**28**

## Presentation 20 -

**15**

## Presentation 15

**16**

## Presentation 16

**17**

## Presentation 17

**18**

## Presentation 18

**19**

## Presentation 19

**24**

## Presentation 24 -

**21**

## Presentation 15 -

**23**

## Presentation 16 -

**25**

## Presentation 17 -

**26**

## Presentation 18 -

**27**

## Presentation 19 -

**30**

## Presentation 22 -

**6**

## Presentation 6 -

**Signal processing, data acquisition**:

**System integration and engineering**:

**Computing**:

**Software and imaging**:

**31**

## Presentation 23 -

**Signal processing, data acquisition**:

**System integration and engineering**:

**Computing**:

**Software and imaging**:

**29**

## Presentation 13 -

**7**

## Presentation 3 -

**9**

## Presentation 8 -

**10**

## Presentation 9 -

**11**

## Presentation 10 -

**12**

## Presentation 11 -

**13**

## Presentation 12 -

**22**

## Presentation 13 -

**20**

## Presentation 14 -

**77**

## Presentation 14 -

**76**

## Presentation 12 -

**5**

## Presentation 3 -

**3**

## Presentation 4 -

**8**

## Presentation 5 -

**84**

## Welcome Welcome (Professor Konstantinos Masselos, Rector of the University of Peloponnese)