# SAFER Grid

Høgskolen i Gjøvik

**Abstract**

We propose the creation of a high performance computing (HPC) environment physcially located at Høgskolen i Gjøvik (HiG) but virtually connected through a Grid environment [1] and therefore providing the computational resources to users nationally and beyond. By combining server environments for off-the-shelf and virtualization workloads as well as a highly innovative high performance solution in the form of the IBM Cell Broadband Engine architecture [5], research problems that would have challenged even well-equipped HPC environments only a short while ago can be tackled efficiently and effectively. Locating the HPC cluster at HiG provides not only important research opportunities for higher education, research institutions, and industry collaborators in the Innland region but also makes synergistic use of the educational facilities at HiG, where a degree program in system administration and management can benefit greatly from a high availability/high performance specimen while ensuring that the environment will be maintained at optimum performance levels. As part of this proposal, several selected current and planned research projects are listed, which demonstrate the utility of the HPC cluster, but are intended mainly for illustration purposes since a considerable number of other research projects will benefit equally or more from this resource.

## 1 Relevance of the Project

Despite the continued advances in computational resources available in desktop and small server systems, there exist a number of categories of research questions in both information security as well as image and media processing that are inaccessible to such small-scale systems but rather require massively parallel processing power and corresponding memory and storage resources. The same also holds true for resources required in modeling and simulation environments applicable in a number of other application areas such as materials science.

Such research infrastructure is also generally not well suited for individual, targeted research proposals either at the national or European level as it tends to amortize itself only over several such research projects and is typically not fundable under such grant proposal guidelines.

By creating a heterogeneous-platform grid component at HiG, a number of local research groups from across the entire scope of the college and the proposed Innland University domain along with projects can be enabled to continue and expand the scope of their research over the course of several years.

These include research on network security and forensics, modeling and simulation, visualization, image and video processing and its applications in both media science and health care and will also enable intensifying not only research and development cooperations and collaborations in the region both among higher education institutions and research establishments but also at the national and Scandinavian level. In addition to targeted research projects, the integration into the Nordic Grid [2] will ensure that the computational resources will be utilized to the best extent possible by accepting workloads from other institutions and, if need be, augmenting its own capacity with additional Grid resources.

## 2 Technical Description

### 2.1 System description

The system proposed is a flexible, high performance computing cluster, that will provide both traditional cluster computing [3, 4], grid computing [1], as well as the flexibillity to allocate parts of the system to running simulations and virtualization of individual computer nodes.

The system is to be based on a hetrogenous architecture, in the form of the two hardware platforms IBM Cell Broadband Engine [5] and x86 systems. The system will be built using high density components, in the form of blade based computer components. The x86 architecture provides an element that is familiar, versatile and flexible, while the Cell Broadband Engine provides a very high performance, low density architecture. Since the Cell architecture is a new area of research regarding high performance installations, it is essential to have an on-site, traditionaly contructed system to assist in the porting and optimization of clustering and grid software and tools.

The system will contain a relatively low number of Cell/B.E. computer nodes, where the theoretical estimate of performance peak is around 5 TFLOPS [6] single precision for 28 processors, 14 computer nodes. To achieve a close-to-par performance on the x86 architecture, an estimate of 200 processor cores is needed, which may be provided by 50 quad-core computer nodes.

The Cell Broadband Engine is a single-source architecture, having IBM as the only manufacturer, and the IBM BladeCenter QS20 as the only high-density computer implementing it. A total of seven BladeCenter QS20 computers may be installed in a single IBM BladeCenter chassis, wich means that two such chassis are needed. X86 blade systems have several sources, a typical configuration is 10 computer blades in one chassis, makin a total of five such chassis are needed.

Storage for the system will be centralized using one or two high density storage cabinets, attached to the computing nodes through either Gigabit Ethernet (iSCSI) or FibreChannel. High density, low cost storage may be provided using modern SATA or SAS based storage cabinets.

## 2.2 Equipment Description

### 2.2.1 Cell Broadband Engine

Sony, Toshiba and IBM has together constructed the Cell Broadband Engine Architecture (Cell/B.E. arch.) [5], a completely new processor architecture. Going into specific details on the architecture is not relevant here, but some explanations on the relevance of this architecture to high performance computing, and justification of cost is of relevance.

The Cell/B.E. is built up by one PowerPC Processing Element (PPE), eight Synergistic Processing Elements (SPE) and the Element Interconnect Bus (EIB). The PPE is a 64 bit PowerPC Generation 5 with Vector/SIMD multimedia extension. This PPC element functions as a general purpose processing unit that can access system management rescources. Hardware rescources defined in the Cell/B.E. architecture are mapped to the real address space seen by the PPE's, giving the PPE the capabillity to access these rescources directly. The SPE's are less complex computational units, and perform no system management functions. They typically process data and initiate data transfers in order to perform their allocated tasks.

The architecture gives very high levels of performance and data throughput [6]. The integrated memory controller provides a peak bandwidth of 25.6GB/s to external memory, the I/O controller provides 25GB/s inbound, 35GB/s outbound, and the EIB supports a peak bandwith og 204.8GB/s intra-chip data transfer among PPSs, SPEs and the memory/IO controllers. The processing elements in total gives a potential performance, as measured with the LINPACK performance testing suite, of 21.03 GFLOPS on double precicion floating point operations, and a peak og 230.4GFLOPS on single precicion. A comparison test with a different architecture, often used in high-performance installations, the Intel Itanium 2 (IA64), performed with the LINPACK testing suite, displayed 14.63 GLFOPS on the given task on the Cell/B.E., while the Itanium2 gave 6.4GFLOPS. Taking into account that the Itanium is regarded as a high-performing processor, and that it's single-precision performance is approximately 2 times its double precision, it may be concluded that the Cell/B.E. performs at 2X the performance of the compared architecture on DP, and at 100X on SP. The numbers listed in this paragraph, regarding performance, are based on actual LINPACK test results, as opposed to marketing estimations.

The IBM Power architecture has for years proven its worth as a high performance computing platform, with its generations Power4 and Power5 (p4, p5) occupying a significant number of top positions at the TOP500 [7] rating list of high performance installations worldwide. There has been done thorough research work on making the Power and InfiniBand [8] architecures a solid foundation to build high performance comuputing systems on. As the Cell/B.E. features a full PowerPC 970 compatible PPE core, a system built on Cell/B.E. will be able to make full leverage of already existing technologies for construction of ditributed computing, i.e. clustering and Grid technologies. Relization of the full potential of the architechture will first be achieved when the new architecture is leveraged, by utilizing the PPE/SPE/MFC/EIB potential. This is an area of research that is completely new, as the architecture is also completely new. This gives the oportunity of not only having a massive computing potential available, but also opening up a fresh area of research. Developing clustering models and toolchains for the Cell/B.E. and performing work on integrating the architecture into existing Grid solutions, e.g. through porting and optimizing the ARC grid controls software for the Cell/B.E. platform, will not only be a solid research area locally, but provide ground work for utilizing the potential of distributed computing on the platform on an international scale.

### 2.2.2 x86 based architecture

While the Cell/B.E. provides a very high performance computing platform, there are many tasks where having available a x86 compatible platform is essential. Running modeling and simulation where previously developed software and models need to be used provide a need for distributed computing capabilities on the x86 platform. Simulation of host/source/target systems for gathering data basis and models in intrusion detection, computer forensics and comunications research are prime candidates for leveraging operating system, hardware and software virtualization. Being able to efficiently evalutate and port clustering and grid toolchains from the commonly used x86 platform to the Cell/B.E. platform will require a legacy system to initiate the porting effort, and for use as a reference platform when performing performance comparisons and systems optimization.

The 64bit extensions to the x86 architecture, in form of the Xeon and AMD64 platforms provide a relatively high performance platform at a low cost. By basing a grid construction on x86 64 blade technology, high performance may be achieved on high density, and at the same time providing high-speed systems interconnects. Connection of such a system to a central storage system will allow for high data rates internally and to storage, both essential in x86 based clustering technologies and at the same time giving flexibillity in the regard that a single node may be allocated a different task, through easily deploy a different operating system and software image to the node.

### 2.2.3 Additional hardware/software requirements

A system designed and built to perform large computational runs and simulations, both locally and through integration into national and multinational grid comuting systems, does naturally have a requirement of high availability. This makes a stable, reliable power supply very important. The total system must include Uninterruptable Power Supply units scaled to the power requirement of the system, with available overhead capacity, and must be built in a redundant configuration.

Operating system software, virtual machine images, and in particular high performance and simulation computing data all require significant amounts of storage space. A system of notable scale does require a dedicated, central storage system that provides a significant storage space.

Controlling nodes are comptuters that are part of a cluster/grid construction, that are normally physically separated from the tight integration of the rest of the system. Each segment of a hetrogenous grid system should have at least one relatively high performing controller node. In the proposed dual architecture system, it will be natural to include one controlling node for each architecture, and separating these physically from the blade based contruction by using single or doule unit rack mountable server computers as controlling nodes.

To allow a maximum available system throughput, and allowing the total system to be as selfcontained as possible, dedicated networking hardware and interconnect should be built into the system. At least two Gigabit Ethernet network connections should be available to any one node in the system, whith room for expansion. The intectonnect hardware chosen for this purpose should be selected for low latency.

All component in the system are available as rackmountable elements. By use of high density constructions, like blade computers/chassis, single unit computers and supporting structures, and high capacity storage arrays, the entire system should be possible to fit into no more than two 42 unit 19 inch rack enclosures.

In general, operating system and clustering/grid software is available either as open source, or through development and/or volume licensing agreements already entered by HiG. However, software for virtualization, and management of storage, virtualized systems and backup is a cost that needs to be accounted for.

# 3 Problem Statements

The following section lists several current projects which will benefit greatly from the added resources available or would make such research projects possible to begin with. While several projects were listed for their close match to the computational resources proposed – particularly the massively parallel computations possible by the Cell environment – the proposed mix of resources will provide excellent flexibility for other types of workloads.

## 3.1 Movie restauration

With the advent of digital media and increasing computational power, restauration of old movies has become an active field of research. The problem of film restauration involves different topics as dust and dirt removal, noise suppression, image stabilisation, scratch removal, flickering correction, dynamic density-correction, frame interpolation, and colour conversion. It has recently been demonstrated [9] that Retinex-based [10, 11] algorithms, such as ACE [12], originally designed for automatic colour correction of still images, can do a good job also for colour correction of strongly degraded movies.

However, such algorithms are inherently computationally expensive, being $O(N^2)$, where $N$ is the number of pixels in the image, per frame, when the processing is performed per frame. For movie restauration, preliminary results indicate that also the temporal variations should be included explicitly in the calculation, turning the complexity into $O(N^3)$. The good point is that the methods applied are trivially paralell by nature in that each pixel is treated independently. With access to a high performance computing cluster, full scale experiments on movie restauration could be performed.

## 3.2 Direct and inverse modelling

At the beginning of the last century, Hadamard [13] formulated the criteria for a well-posed mathematical model of a physical phenomenon:

1. A solution exists.
2. The solution is unique.
3. The solution depends continuously on the data, in some reasonable topology.

Direct mathematical modelling satisfies these criteria. The fundamental conservation equations and the constitutive equations together constitute a set of laws, $L_i$, whereas the boundary and initial conditions along with the parameters entering the constitutive equations constitute a collection of input data, $c_i$, in the model. The direct model is then in full correspondence with the deductive-nomological (or D-N) model according to Hempel and Oppenheim [14]:

$$\text{if } c_i \text{ and } L_i \text{ then } X_i,$$

where $X_i$ are the events to be described, i.e., the outcome of the model.

Inverse problems arise when we study the outcome $X_i$ of a physical process governed by given laws $L_i$, and want to find some of the input data $c_i$, e.g., finding some parameters of a constitutive law, determine a boundary or initial condition, or simply optimise the process. We are then following the opposite scheme,

given $X_i$ and $L_i$, what must $c_i$ have been?

This is the general scheme of an inverse problem, and, in general, it is not well-posed according to the definition by Hadamard [13]. A modern introduction to the topic of inverse problems is given by, e.g., Tarantola [15], and the specific case where the direct model consists in partial differential equation that must be solved numerically, i.e., *inverse modelling* is thoroughly treated by Isakov [16].

Solving the direct problems is an important task in materials technology industrial research. For solving the direct problems in continuum mechanics, standard numerical methods for partial differential equations are commonly used such as the finite element method (FEM) [17], either through commercial codes (e.g., ABAQUS) or through custom-built codes. Solving inverse problems built upon such models involves among other things solving a number of parallel direct models. In other words, the problem is embarrassingly parallel and can thus easily be given a parallel implementation [18]. Since the solution of the direct model is itself computationally expensive, a parallel implementation of the inverse modelling problem is well suited for clustering [3, 4] or grid [19] technologies. In this way, more exact solutions can be found, and larger regions of the parameter spaces can be exploited.

### 3.3 Critical Infrastructure Modeling and SImulation

The modeling and simulation of critical infrastructures (which includes a number of sectors including e.g. telecommunications, energy, transportation, financial services, health care, and government services) is a vital research challenge to better understand the vulnerabilities of society as a whole and to devise possible mitigation strategies. This is the case not only with regard to critical information infrastructure but also the interfaces between information and physical infrastructures as exemplified by the tight integration of control systems and IT services in any of the core infrastructure sectors at the regional, national, and in some cases even international scale. Threats to this infrastructure can come from diverse sources including natural disasters, human error, but also inimical nation-states and sub-state actors while the outcome of damages to the critical infrastructure can reach from loss of economic output via loss of life or limb to threats to national sovereignty. It is therefore essential to understand the interactions and vulnerabilities of the increasingly interconnected infrastructure not only from the perspective of component reliability. Rather, it is dynamic effects and perturbations (e.g. as may be the case for targeted attacks) that are of particular interest which may spread and resonate across longer time periods and (logical and geographical) extents as indicated by current research.

However, both the size and predictive accuracy of current models must be increased significantly to ensure that relevant research questions can be answered. In addition, algorithmic optimization (both numerical and particularly combinatorial) approaches can identify vulnerable infrastructure areas and can also be identified effective and efficient ways of improving overall robustness and resilience.

This requires models which go beyond currently existing ones in terms of their precision for individual infrastructure sectors that may approach existing sector-specific models together with their interactions along with sufficiently high temporal resolutions. Moreover, since many elements of critical infrastructures still operate with humans in the loop, such models may also require agent-based simulations as part of the overall research approach. This in turn requires the use of Grid-level massively parallel computation and memory resources. In addition, however, the design and development of models themselves and algorithms operating over these can benefit from intemediate-scale systems for faster exploratory research before committing to extensive simulation runs.

Research at HiG (in cooperation with national and international partners) that is ongoing, submitted to NFR, and also planned submissions to NFR and at least one research proposal under preparation for the joint ICT/Security call of the 7th Framework Program of the European Commission will greatly benefit from the availability of a flexible high performance compute cluster together with access to tap into the Nordic Grid for large-scale computations which will also be linked with related visualization efforts both locally and within the proposed EU FP7 project consortium.

## 3.4 MANET Security Network Simulation

Mobile ad hoc networks and sensor networks have become the focus of intense research over the course of the last decade, with security of such environments also playing a major role.

Research thus far has, however, concentrated mainly on small-scale, idealized network environments and on sub-problems such as routing security which do not take into account a number of environmental parameters and information which should be used to assess and improve the security posture of the network. At the same time, the quality of the simulations has generally been hampered by the need to omit realistic sub-models for radio signal propagation particularly in built environments, advanced mobility models, and models encompassing longer time-spans (minutes to hours as opposed to the seconds to minutes currently being simulated).

These parameters hide significant research potential and surprising results and are key to investigating the use of MANET and sensor networks regarding overall information security – as opposed to isolated aspects such as provable or probabilistic route security or the detection of certain types of attacks – and robustness in less beningn application areas such as law enforcement, emergency responders, civil defense, and military areas.

While one approach to address this problem, e.g. taken by some military organizations, is to develop hybrid approaches in which larger-scale but coarse models are coupled with small-scale experimental systems, the continuation of present research into improving the simulation capabilities of network simulation and in the process investigating issues such as efficient distributed computations for trust infrastructures or other infrastructure services required in MANET network holds greater promise owing to its flexibility. As significant portions of such simulations and models can be parallelized and distributed, a Grid-based computational model approach is the most appropriate one to address the significant requirements arising for this type of research and will not only enable the continuation of current research but also the deepening of existing research collaboration with leading academic and government laboratories.

## 3.5 Network Security and Forensics

Both the validation of network simulation results as described in section 3.4 and research not amenable to simulation can benefit greatly from actual deployment of new algorithms and mechanisms onto intermediate-sized networks. This allows the analysis of effectiveness and performance characteristics as well research into computer and network forensics under circumstances that are both controlled and sufficiently realistic to yield relevant results.

A key problem faced by such research thus far has been the significant effort required to configure such test-bed networks to ensure that results are reproducible and measurements are not affected by configuration artifacts. However, the availability of hardware-supported virtualization on commodity platforms allows not only the multiplexing of a larger number of virtual systems into a virtual network but also the creation of precisely circumscribed configurations which can then be used to set up reproducible experimental environments.

Examples of such research includes forensic research on virus and worm behavior and propagation and communication characteristics which may not be possible in honeypot or honeynet environments as these lack the critical element of reproducibility and inspectability of virtualized platforms. The behavior and interactions of higher-level protocols and user activity is also relatively poorly understood but is critical for the functioning of security and privacy protocols. Once again in this case, the availability of flexible, virtual, and well-characterized experimental networks and system configurations are critical for obtaining statistically significant results that are difficult to obtain in regular experimental environments.

Finally, as has been documented earlier, the same network can also play a vital role beyond its applications in research in network security and administration education as well and will provide synergistic effects for education at B.Sc., M.Sc., and Ph.D. levels both for managing and securing networks and also for the unique challenges posed by high performance and high availability grid computing environments.

# 4 Educational, national and international relevance

## 4.1 Education, IT management

A cluster-configuration can give useful experiences in the area of IT-management, in particular view that recent years has shown a significant focus on economy of scale with advantages clustering and 24/7 services based on failover solutions.

A set of nodes in a computing cluster provides a good angle at learding how large-scale solutions are build, configured and maintained. A large share of IT-systems today are outsourced services in large datacenters. For a person that aims towards a career in IT administration/management, it is highly relevant with a position with a larger IT service partner. In running large scale solutions, standardization, automated maintenance and servicing are of utmost importance. It is therefore essential that students in IT management education are familiar with the methods and principles behind this.

There are also trends in the industry that indicates that dynamic, mobile, modular datacentres are going to become more prominent. En example of this idea is the artchitectural ideas of the infrastructure used by e.g. Google. A small-to-medium x86 based clustering solution is going to need the same solutions as a large scale data center, and will provide students a view and experience with how such a centre needs to operate.

Another factor is that the computational rescources in a cluster may be harnessed to deploy virtual machines (appliances) for use in education. Each student can be provided an appliance that may be set up and configured according to need. Such an appliance will only run when needed, and will as such not tie up rescources in the permanent manner that dedicated workstations and terminal servers do. At the end of semester, this rescouce can also easily be deallocated and used for other purposes.

## 4.2 Connection to NorduGrid

NorduGrid [2] is a Grid Research and Development collaboration which develops, maintains and supports the so-called Advance Resource Connector (ARC), a free Grid middleware. The NorduGrid activity started out with a project called the "Nordic Testbed for Wide Area Computing and Data Handling". That project was launched in May 2001, aiming to build a Grid infrastructure suitable for production-level research tasks. The testbed was set up and remain in continuous operation and development since August 2002. Currently, there are more than 50 sites from 12 different countries connected to NorduGrid, with more than 6000 CPUs available for computational purposes. The main application area of NorduGrid is data processing for high-energy physics, in particular for experiments at the new Large Hadron Collider accelerator at the research laboratory CERN outside Geneva, Switzerland.

A computing cluster at HiG would potentially enable HiG to become part of the NorduGrid collaboration, as the cluster could become part of the Nordic Testbed. In addition, research work on security issues relevant to Grid computing carried out at HiG would be highly welcome inside the NorduGrid collaboration.

In return, a connection to NorduGrid implies access to resources far beyond that of the local cluster itself. This would be important for computer-intensive research at HiG, but it would also open up unprecedented opportunities for computer-intensive research projects within material technology carried out in collaboration with the local Norwegian Centre of Expertise (NCE), located at the industrial park in Raufoss a few kilometres away from HiG. Access to such computing resources is highly desired inside NCE Raufoss and would further increase HiG's attractivity as a serious research partner.

### 4.3 WebDeal / IBM

A significant partner in the aquisition of parts of the proposed system, is WebDeal and IBM. WebDeal is a national IBM partner, that has a long running research cooperation with IBM regarding the Power platform, participation in the World Community Grid, and a long history of collaboration with HiG in respect of educational and research projects. WebDeal as a local partner has a significant interest in an installation of IBM Cell/B.E. clustering technology. Interests here include research on maintaining and running a large scale system based on this architecture, integration of the architecture in Grid contructions, and performance measuring of the architecture. It is viewed as a highly relevant performance measure to apply real scientific problems and evaluating the performance result in comparison with a more standard x86 architecture. This performance measurement idea is deemed far more telling than what has been done till now; running benchmarking systems and comparing results of those.

### 4.4 Frozt Film

Frozt film is a multi-national entity that focuses on digital video production, processing and distribution. Their primary business location is local in respect to HiG, as their base is at Gjøvik. Frozt film has expressed that a high performance computer installation would be of great application to them for use in image processing and content rendering, but is outside of direct availabillity to them. Frozt is very interrested in access to such an installation locally, and thus has a direct interest in the installation of such at HiG.

## 5 Estimated budget

The cost outlined here is an estimate based on a rough approximation of cost. As parts of the system is single-source, and the equipment being new technology, an absolute, precise cost is currently difficult. Because of this, an approach where a realistic estimate of cost and need for funds is provided, and these numbers will be used to negotiate a precise offer from the suppliers once funding is guaranteed. As the proposed installation is of some scale, this is the approach that is preferred by the technology suppliers. The numbers for standard components, like networking, UPS and storage, are based on previous experience, but are also to be seen as estimates.

| | |
|---|---:|
| IBM BladeCenter with Cell blades | 2 000 000 |
| x86 blade nodes with chassis | 1 200 000 |
| Networking equipment | 100 000 |
| Controlling nodes | 100 000 |
| Storage | 300 000 |
| UPS | 100 000 |
| Software costs | 150 000 |
| Other hardware | 50 000 |
| Estimated sum | 4 000 000 |

# 6 Strategic Perspective

## 6.1 Benefits to Society

Most of the benefits that will emerge from the proposed system will be derived from the individual research projects utilizing this shared resource and are therefore difficult to quantify immediately. However, it is already evident from the specimen research projects outlined in this proposal that society will benefit e.g. from improvements in the preservation of cultural heritage through advanced restoration techniques for movie materials, or enhanced predictive capabilities of critical infrastructure models and simulations which may help in limiting environmental damages of natural and man-made disasters as well as reducing the risk of life and limb because of incidents involving critical infrastructure. Finally, by providing regional and national industry with enhanced modeling and simulation capabilities, the competitiveness of these industries is strengthened.

Beyond these benefits, there are also indirect benefits to be derived from enhanced regional, national, and international academic and scientific collaborations which will not only lead to advancements in knowledge but must also be considered a net benefit in and of themselves.

## 6.2 Environmental Impact

The proposed system is not anticipated to have an environmental impact on its own. It will, however, provide the opportunity to save substantial amounts of energy by allowing the migration of workloads onto highly efficient server systems that would otherwise require the use of servers or server clusters exhibiting significantly lower energy efficiency.

Moreover, by integrating the resources into a Grid environment, these can be made available to other researchers both in the area and particularly at the national level and beyond, which can substantially reduce the environmental impact of travel costs normally required when visiting high performance computing facilities.

## 6.3 Ethical Issues

## 6.4 Equal Opportunity

All project partners are committed to the principles of diversity and equal opportunity. In particular, the recruitment of the project staff will be in accordance with current regulations, i.e. recruitment will be in accordance with the "*Handlingsplan for likestilling i FoU-sektoren*".

# 7 Communication and Exploitation of Results

## 7.1 Communication with Users

## 7.2 Dissemination and Exploitation Plan

# References

[1] A.Wäänänen et al. An Overview of an Architecture Proposal for a High Energy Physics Grid. Applied Parallel Computing. Advanced Scientific Computing: 6th International Conference, PARA 2002, Espoo, Finland, June 15-18, 2002. Proceedings *Lecture Notes in Computer Science* vol 2367/2002 Springer-Verlag Berlin Heidelberg 2002.

[2] http://www.nordugrid.no

[3] Rajkumar Buyya, editor. *High Performance Cluster Computing: Architectures and Systems*, volume 1. Prentice Hall, NJ, USA, 1999.

[4] Rajkumar Buyya, editor. *High Performance Cluster Computing: Programming and Applications*, volume 2. Prentice Hall, NJ, USA, 1999.

[5] IBM, October 2006 Cell Broadband Engine Archtecture, 1.01

[6] Thomas Chen, Ram Raghavan, Jason Dale, Eiji Iwata. Cell Broadband Engine Architecture and its first implementation: A performance view

[7] http://www.top500.org

[8] http://www.infinibandta.org

[9] Alessandro Rizzi, Davide Gadia, and Daniele Marini. Analysis of tri-stimulus interdifference and contextual color correction. *Journal of Electronic Imaging*, 2006. Accepted.

[10] Edwin H. Land and John J. McCann. Lightness and retinex theory. *Journal of the Optical Society of America*, 61(1):1–11, January 1971.

[11] Edwin H. Land. The retinex theory of color vision. *Scientific American*, 237:108–128, 1977.

[12] Alessandro Rizzi, Carlo Gatta, and Daniele Marini. A new algorithm for unsupervised global and local color correction. *Pattern Recognition Letters*, 24:1663–1677, 2003.

[13] J. Hadamard. *Lectures on the Cauchy problem in linear partial differential equations*. Yale University Press, New Haven, 1923.

[14] Carl G. Hempel and Paul Oppenheim. Studies in the logic of explanation. *Philosophy of Science*, 15:135–175, 1948.

[15] Albert Tarantola. *Inverse problem theory and methods for model parameter estimation*. Society for industrial and applied mathematics, 2005.

[16] Victor Isakov. *Inverse Problems for Partial Differential Equations*, volume 127 of *Applied Mathematical Sciences*. Springer-Verlag, 2nd edition, 2006.

[17] Kenneth Eriksson, Claes Johnson, Donald J. Estep, and Peter Hansbo. *Computational Differential Equations*. Cambridge University Press, 2nd revised edition edition, 1996.

[18] Hans P. Langtangen and Aslak Tveito. *Advanced Topics in Computational Partial Differential Equations: Numerical Methods and Diffpack Programming*. Lecture Notes in Computational Science & Engineering. Springer-Verlag Berlin and Heidelberg GmbH & Co., 2003.

[19] M. Ellert, A. Konstantinov, B. Kónya, O. Smirnova, and A. Wäänänen. The nordugrid project: using globus toolkit for building grid infrastructure. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 502(2-3):407–410, April 2003.