Jornada de Seguimiento de Proyectos, 2005
Programa Nacional de Tecnologías Informáticas

Improving the performance and services offered by PC clusters. Development of distributed multimedia applications TIC2003-08154-C06

José Duato [*]
Universidad Politécnica de Valencia

Francisco Quiles [†]
Universidad de Castilla-La Mancha

José Manuel García[‡]
Universidad de Murcia

Juan Manuel Orduña[§]
Universitat de Valencia

Francisco Mora[¶]
Universidad Politécnica de Valencia

José Manuel Claver[‖]
Universitat Jaume I

Abstract

PC Clusters have become the most attractive platform to implement supercomputers and high-performance servers. As a consequence, it is necessary to improve these platforms in order to provide higher performance and scalability, and better services, such as availability, reliability, quality of service, low power, etc.

In this project, we propose to improve most of the hardware components of a PC cluster. We will also improve the cluster system software. In particular, we propose to improve processor performance by using superpipelining techniques and by reducing power consumption. I/O system performance will be improved by using a switch fabric with point-to-point links. SMP scalability will be improved with a new memory coherence scheme. With respect to the interconnection network, we propose improving not only its performance (by using adaptive routing and load balancing techniques) but also its scalability (by using congestion management techniques), availability (by using fault-tolerance techniques) and quality of service (by including some hardware/software support in the routers and network interface controllers). We will also develop scheduling techniques that will take into account communication costs. Regarding Internet access, we will enhance IP routers, improving their scalability and bandwidth by means of new congestion control mechanisms and HOL blocking avoidance techniques. We will also research on wireless networks, developing new robust coding schemes based on new UEP techniques and applying them to the transmission of MPEG-4 video. Additionally, we will research on distributed virtual environments, developing new

[*]Email: jduato@disca.upv.es
[†]Email: paco@info-ab.uclm.es
[‡]Email: jmgarcia@ditec.um.es
[§]Email: Juan.Orduna@uv.es
[¶]Email: fjmora@eln.upv.es
[‖]Email: claver@inf.uji.es

partitioning methods capable of improving the efficiency and scalability of these systems. We will also develop some distributed applications focusing on those that require some specific hardware to allow real-time execution. In particular, we will develop quantization and coding procedures and their hardware implementation on FPGAs. Finally, we will improve the synthesis of FPGA circuits from high-level specifications.
Keywords: PC clusters, interconnection networks (adaptive routing, congestion management, fault tolerance, network reconfiguration, QoS support, power consumption), wireless networks, shared-memory multiprocessors, circuit synthesis, video transmission, load scheduling, distributed virtual environments.

1 Project objectives

This research project has two primary goals: Improving the performance and services offered by PC clusters, and developing distributed multimedia applications. Each one of these objectives can be split into several subobjectives, organized into subprojects:

- Objective 1.1. Improving the performance and scalability of the interconnection network.
- Objective 1.2. Improving the performance of current interconnection technologies for system area networks (SANs).
- Objective 1.3. Improving the performance of Internet routers.
- Objective 2.1. Improving the quality of service and availability provided by the interconnection network to the applications.
- Objective 2.2. Design of switch architectures optimized for the execution of applications that broadcast multimedia traffic.
- Objective 2.3. Improving the performance of the transmission of MPEG-4 encoded video through wireless local area networks and development of new error-resilient encoding techniques.
- Objective 2.4. Analysis and dynamic parameter tuning for the medium access control layer in the IEEE 802.11 a/b WLAN.
- Objective 3.1. Improving performance and developing power consumption reduction techniques for the cluster nodes.
- Objective 3.2. Improving the services offered by PC clusters.
- Objective 3.3. Efficient development of applications for high-performance platforms.
- Objective 4.1. Improving the scheduling techniques for parallel systems.
- Objective 4.2. Improving the performance and scalability of distributed virtual systems.
- Objective 5.1. Development of a quantizer and real-time video encoder amenable for implementation in a reconfigurable VLSI architecture.
- Objective 6.1. Automatic FPGA programming from high-level specifications.

The time chronogram has been omitted since it occupies 7 pages.

2 Project results

2.1 Research results

This section summarizes the activities developed from the beginning of the project, grouped by objectives and topics within each objective. For each topic we highlight the most relevant scientific and technological results that have been obtained.

2.1.1 Objective 1.1.

A) Reduction of HOL-blocking (congestion control). We have designed a new technique, named RECN (Regional Explicit Congestion Notification), that completely eliminates the HOL blocking produced by congestion trees. When congestion is detected, queues are dynamically allocated to store the packets flowing through the congested point, thus isolating congested traffic and avoiding HOL blocking. This mechanism is able to handle congestion trees that grow from root to leaves and from leaves to root, overlapping congestion trees, and tree merging. RECN only requires up to 8 additional queues per port, independently of network size. RECN is very fast and, to the best of our knowledge, the only truly scalable congestion management technique proposed to date. RECN has been patented in the USA jointly with Xyratex. We also developed another technique to reduce HOL blocking, named DBBM (Destination Based Buffer Management), which is very simple yet effective. It is based on assigning groups of packet destinations to each of the queues, so that each queue can only produce HOL blocking among the packets it contains. The main benefit of this technique is its low cost. However, it does not completely eliminate HOL blocking.

B) Development of adaptive routing techniques. Solutions to out-of-order packet delivery. We have continued the work aimed at improving SAN performance by using adaptive routing strategies. The proposed strategies have been evaluated under different scenarios. One of the main drawbacks of adaptive routing is that it can not guarantee in-order packet delivery, feature that is required by many applications. During the last year, we have addressed the out-of-order packet delivery problem. To this end, we have analyzed, under different scenarios, the number of packets arriving out-of-order at destination, also evaluating the number of resources required for guaranteeing in-order packet delivery. Additionally, we have proposed two effective approaches to tackle the problem. One approach is based on allocating a very limited number of reordering buffers at the destination nodes together with certain end-to-end flow control mechanism. Another approach is based on limiting the packet injection at the source nodes, reordering packets at destination if needed. We have shown that it is possible to guarantee in-order packet delivery when applying adaptive routing by using a small number of additional resources and without significantly penalizing network performance.

C) Balancing techniques based on traffic distribution. We have proposed a cost-effective evaluation methodology for traffic balancing algorithms that is based on the use of genetic algorithms. This methodology obtains a near optimal routing path distribution from the applied routing algorithm, allowing us to asses the effectiveness of the traffic balancing algorithms commonly used in PC clusters and SANs.

D) Fault-tolerant routing for direct networks based on the use of intermediate nodes. We have developed a new methodology that, in case of failure(s), selects an intermediate node and routes every packet

traveling through the faulty region in two phases: from the source to the intermediate node and from there to the destination. This methodology is able to tolerate a relatively high number of failures without degrading performance in the absence of failures, and degrading it minimally when failures occur in the interconnection network. In addition, these features are achieved, unlike for previous proposals, without disabling healthy nodes, without using too many hardware resources, and introducing no penalty for most packets (for example, an increase in latency) when they are routed in a network with failures.

E) New strategy for interval routing. This strategy has been developed for interconnection networks built from commercial switches and has been optimized for the more widely used network topologies (tori and meshes). By optimized we mean that it is scalable from several points of view: (i) the memory space required to store the routing tables does not linearly depend on network size, (ii) the hardware required to apply the routing algorithm is not complex and does not depend on network size, and (iii) the time required to perform a routing operation does not depend on network size.

F) Power consumption reduction in interconnection networks. We have investigated two different schemes: The first proposal adjusts the link width based on the current load in the network. In this case, no link becomes completely disconnected. Therefore, the same routing algorithm can be used independently of whether the power saving mechanism is active or not, which simplifies the router design. This strategy is able to achieve very large reductions in power consumption. The main disadvantage is the increase in the latency when the network is lightly loaded. The second proposal reduces power consumption in interconnection networks based on the use of a regular topology and high-degree switches with channels consisting of multiple serial links running in parallel. This approach is based on the selective connection and disconnection of the high-speed serial links that compose the channels, depending on the current amount of traffic in the network. The behavior of the mechanism depends on the election of two thresholds. The design restrictions regarding the election of the thresholds and the relationship between these parameters and the behavior of the network in terms of latency and power consumption have been carefully analyzed. With a careful election of the thresholds, the system can be tuned to obtain different levels of power saving with a moderate increase in latency.

2.1.2 Objective 1.2.

A) Strategies to compute routing tables for SANs. We have proposed an effective methodology to compute up*/down* routing tables in SANs. This methodology shows that just assigning directions "up" or "down" to the links in an effective way is able to significantly increase the performance achieved by up*/down* routing (throughput increases up to a factor of 2.8).

B) Deadlock-free minimal routing techniques for SANs. We have proposed a new generic routing algorithm for SANs that is able to provide minimal routing by using a number of resources smaller than that required by other approaches. The new routing scheme, developed in collaboration with the Simula Research Laboratory in Oslo, combines the idea of classifying the set of minimal paths and assigning each class to a different virtual network, with the idea of enabling transitions between virtual networks.

C) Fault-tolerant routing techniques for InfiniBand. We have proposed an effective and scalable fault-tolerant routing methodology for InfiniBand under dynamic fault model, which takes advantage of the

automatic path migration (APM) mechanism provided by this network technology. The methodology is able to provide the maximum number of disjoint paths allowed by the topology. Deadlock-free disjoint paths are computed by allowing transitions between virtual channels in an ordered way. Recent results focused on providing fault tolerance by combining routing through intermediate nodes and partial network reconfiguration techniques.

D) Congestion reduction techniques for SANs. We have proposed an effective congestion management strategy for InfiniBand networks based on marking and validation of in-transit packets. This strategy allow us to thoroughly identify the traffic flows giving rise to the congestion and to measure the congestion degree in the network in order to apply effective corrective actions.

2.1.3 Objective 1.3.

See Section 2.2.

2.1.4 Objective 2.1.

A) Switch-fabric design of QoS-aware IP routers. See Section 2.2

B) QoS support in InfiniBand. We have evaluated the performance of InfiniBand via simulation under realistic traffic loads and by changing various system parameters. Specifically, we have performed a study on the buffer space required to meet the QoS needs as specified at connection time by the applications. From our results, we have found out that a buffer space of one can fulfill the application requirements. We have also analyzed the impact of the packet size over the QoS as perceived by the end application. In particular, we have found out that the use of small packet sizes implies a higher level of overhead.

Moreover, we have also developed a complete theory to verify the correctness of our proposal. Our theory has been developed from our previous work on a QoS-aware traffic classification. We propose to configure the InfiniBand mechanisms to guarantee the QoS requirements of a wide variety of applications. Currently, this proposal is being implemented into a prototype making use of a commercial InfiniBand system.

C) Improving the availability of the interconnection network. Under this activity, our main results are the following. First, we designed and evaluated a partially implicit routing algorithm (PIRa), that quickly (25% of the time required by previous alternatives) obtains a valid set of routes. It is based on the utilization of default ports in the upward segment of every up*/down* route.

Next, we designed a hybrid mechanism that combines PIRa with other traditional FERa (fully explicit routing algorithms) to mix their respective advantages, speed of convergence and network performance.

Finally, we defined two deadlock-free dynamic routing table distribution processes based on the deactivation of break ports and break dependencies, respectively. Both mechanisms outstandingly improve the results provided by previous proposals. In particular, the behavior of the second one is quite similar to the one expected for a hypothetic ideal dynamic distribution process.

D) QoS support for source routing SANs. We have proposed several methods to use these mechanisms to provide QoS based on bandwidth and latency requirements. We have evaluated the behavior of our proposals using CBR and VBR traffic. We have compared the VC arbitration table scheduler with our WFQ Credit Aware implementation of the MinBW, using CBR and VBR traffic. The results obtained have shown that both methods are able to provide a differentiated treatment to the different service classes. Results have also shown that the VC arbitration table is the only specified mechanism that can be used to provide both bandwidth and latency guarantees. However, the WFQ Credit Aware method obtains very good results for both bandwidth and latency, but it is unable to provide guarantees. The results obtained have shown that the injection flows must be controlled in order not to exceed the reserved bandwidth limits. We have found out that even if the flows are not controlled, the bandwidth requirements still can be guaranteed, but not the latency requirements.

2.1.5 Objective 2.2.

The main objective of this activity has been the design of a Multimedia Router (MMR) architecture: a single-chip router to be used in local area networks and cluster environments. The architecture has been designed to efficiently handle CBR and VBR multimedia traffic as well as best-effort traffic. An in-depth performance evaluation of various proposed link and switch scheduling algorithms has been carried out via simulation. As part of this effort, we have developed a simulation tool.

During this year, we have been enhancing the proposed architecture. In particular, we have proposed a novel link scheduling algorithm: Temperature-IABP (TIABP). Our results show that the use of this algorithm can significantly reduce the hardware complexity of the switch. An FPGA implementation of the proposed architecture clearly shows that it is feasible. Our results also show the effectiveness of the proposed architecture in meeting the QoS requirements of the various applications.

2.1.6 Objective 2.3.

Two main results have come out of this research activity: 1) the optimization of robust coding techniques UEP to be used by DCT-based coders, such as, MPEG-2 and MPEG-4, and 2) an analysis of the impact of the handover procedures over the video quality services. Throughout the first phase of the project, we have focused our efforts on the MPEG-4 video coding standard. Due to new developments in the area of video coding, we have recently undertaken the design of fast transcoding techniques. In particular, we are designing fast prediction techniques to be incorporated into the MPEG-2/H.264 transcoding process. Our choice has been based bearing in mind that the H.264 standard will become the most widely used video coding standard in the near future. It is also well known that the MPEG-2 standard has been widely used, i.e., there exists a large quantity of MPEG-2 encoded video material. Furthermore, the results of this research task should enable the design of transcoding techniques that in turn will enable the deployment of adaptive video services over WLANs.

2.1.7 Objective 2.4.

Under this activity, we have first carried out an in-depth analysis of the most up-to date QoS mechanisms for IEEE 802.11 WLANs. We have then developed various simulators and carried out an extensive set of simulations. Throughout this study, we have considered a WLAN cell working in the infrastructure mode supporting MPEG-4-based video applications. The various QoS mechanisms under study are enabled by modifying the system parameters: IFS, CWmin, CWmax. We have also

undertaken the study of the EDCA (802.11e) mechanism. EDCA is called to become the IEEE standard for QoS provisioning in IEEE 802.11 WLANs. Recently, we have designed a QoS framework fully compatible with EDCA. This mechanism is fully compatible with the DCF mechanism of the IEEE 802.11 standard. This feature ensures the interoperability of DCF-based stations with the EDCA mechanism: a key feature given the large number of DCF-based stations already deployed.

2.1.8 Objective 3.1.

A) Performance improvement and power consumption reduction for processors. We have performed an analysis of how superpipelining affects the energy consumption in high-performance processors. We evaluated pipeline depths from 5 to 42 stages, seeing how power consumption grows, mainly due to the increasing speculation. We made a proposal that reduces the power consumption by 16% for 42 pipeline stages.

Also, we have looked at embedded processors power consumption. We have proposed an instruction fetch unit that makes a more efficient use of the instruction cache. This mechanism, called "Fetch Mask Predictor", identifies which instructions of each line are actually going to be used and performs a partial access. This way, we obtain a reduction of 44% in the power consumption of the instruction cache, contributing to improve the reliability, stability and lifetime of processors and batteries.

Finally, we have looked at power consumption for high-performance processors. In this case, peak currents and dissipated power can cause many faults and even damage the processor. We have generalized the previous proposal to apply it to high-performance processors. Again, doing partial accesses to instruction cache lines allows a significant power reduction without compromising performance.

B) Design of a new generation of Journaling File systems. We have added extensions to DualFS and studied its performance with large files. Using extensions has allowed us to reduce the space used by metadata, specially for large files.

To improve the performance of directory operations we have implemented them using B+ trees. To maintain the transactionality when using B+ trees, we have designed a mechanism called "metafiles" that will allow us to simplify DualFS implementation and add new features in the following years.

Also, we have tuned our current implementation determining the best default values for a number of parameters.

C) Scalable shared-memory multiprocessors. We have adapted our simulation tools to simulate high-performance SMP architectures, cc-NUMA architectures and CMP architectures. We are also studying how to connect the SIMICS simulator to our current simulation tools to be able to simulate commercial workloads.

Also, we are currently developing a scalable directory-based memory architecture and comparing it to high-performance SMP systems.

2.1.9 Objective 3.2.

A) User-transparent resource management for networks of workstations. We plan to develop an efficient resource management system, develop low-overhead communication protocols and develop fault-tolerance support.

We have already developed an environment for running applications based on MPI using novel interconnection network features and a new methodology for cluster software development that enables

better stability and predictability. Our methodology will ease the development of debugging schemes and fault-tolerance features.

B) Development of user management techniques for wireless environments. We are currently analyzing available tools for access control, designing techniques for access control management in structured environments and designing techniques for building ad-hoc networks.

We have designed a first prototype which brings a solution for the integration of authentication and authorization systems in a concrete context of access control, using SAML and XACML standards for information representation.

C) Integration of a component model (CORBA-LC) to the Grid (Globus). We are studying the non-functional requirements for Grid environments, developing a requirements document for component containers specifying the services offered, and implementing and testing the CORBA-LC component container model.

The work on this activity has focused on implementing a code generator for components, adapting the generator to generate Globus Web Services skeletons that offer the services provided by CORBA-LC components and implementing a proxy generator that allows wrapping Globus Web Services as CORBA-LC components.

We have been working on integrating the CORBA-LC component model and the Globus Grid, specially studying the meta-information facilities provided by Globus.

We have developed mechanisms for integrating version 2 and version 3 of the Globus platform, which are incompatible. These mechanisms ease the migration between versions.

Finally, we have proposed a policy-based management architecture (PBM) which aims at providing a flexible auto-configuration method for Grid nodes.

2.1.10 Objective 3.3.

A) Optimization of the compression and decompression processes for medical video using the 3D wavelet transform. We have studied and applied multithreading technology to video compression using the 3D wavelet transform. We have studied the behavior of general-purpose processors with SMT capabilities (Intel Pentium IV Xeon with Hyper-Threading) running the video encoder. We have proposed several parallelization techniques using different work decompositions. We have obtained speedups as high as 1.34 using Hyper-Threading and a functional decomposition. Additionally, we have implemented this parallelization scheme using both OpenMP and Pthreads to compare both implementations with respect to performance, ease of implementation, and maintainability of the resulting programs.

We have performed several studies comparing manual vectorization using SIMD instructions and automatic vectorization using modern compilers. We have found that automatic vectorization is not good enough yet. After studying the different SIMD instruction sets (Intel's MMX, SSE and SSE2; AMD's 3D-NOW; PowerPC's AltiVec; Sun's VIS; HP's MAX; and MIPS's MDMX) we have concluded that the available SIMD instruction sets allow the execution of the 3D wavelet transform in real time and no new instructions need to be proposed to achieve that goal.

B) Design, implementation and acceleration of image processing algorithms for artificial vision. Our work has focused on developing the basic infrastructure for the implementation and optimization of algorithms for image processing and artificial vision using conventional processors with multimedia and vector processing extensions. In particular, we have reached a fairly advanced development state for a

general purpose library for these purposes called SVNVISION which is based on Intel IPP (Integrated Performance Primitives) and achieves high speed in basic tasks for image processing.

Finally, we are preparing an extensive work about auto-calibration for optical sensors for autonomous agents using odometry.

2.1.11 Objective 4.1.

A) Development of a joint and synchronized simulation tool, capable of simulating both the execution of processes and the interconnection network. We have designed a simulation tool capable of jointly simulating the execution of processes as well as the interconnection network. Also, we started the implementation of such a tool. Currently, we are at the stage of testing the alpha version of this tool.

2.1.12 Objective 4.2.

A) Development of heuristic techniques for solving the partitioning problem in Distributed Virtual Environments. We have implemented several different heuristic search methods (including constructive, evolutionary and bio-inspired heuristics) that provides better partitions (assignments of clients to the servers in the system) while requiring shorter execution times than the methods currently proposed in the literature.

B) Characterization of Distributed Virtual Environment (DVE) systems. Surprisingly, none of the partitioning methods proposed in the literature had been correlated with system performance. We have correlated the partitioning method that is currently considered to provide the best results in terms of system performance. Since the correlation results have shown a total absence of correlation, we have studied the performance of DVE systems as the two main parameters in the partitioning method (average message latency and average CPU utilization of the servers in the system) vary. We have shown that, in order to design efficient and scalable DVE systems, the partitioning method should be targeted at avoiding that any of the servers in the system reaches 100% of CPU utilization, regardless of the message latency.

C) Development of dynamic, efficient partitioning methods for DVE systems. Taking into account the results obtained in topic B), we have developed a new partitioning method, based on a dynamic load balancing technique, that actually improves the DVE system throughput without significantly affecting system latency.

D) New goals not defined in the initial proposal. Due to some difficulties described below, we have led our efforts to propose and fulfill new goals. In particular, we have designed and implemented a new partitioning method, based on heuristic search, that not only maximizes system throughput but it also provides Quality of Service to the maximum number of clients as possible. This method starts from the method described in topic C) and uses the remaining bandwidth (if any) to reduce the latency provided to avatars to a given threshold value.

2.1.13 Objective 5.1.

A) Study of the needs for the image compression system. We have done a deep study of the different techniques—from quantization and coding of images and video based on the wavelet transform—, to

determine the specifications of the electronic system that we have to design.

B) Development of a test bench. We have developed a test bench for verifying the image compression system. The tool VERA has been used for the verification of the digital electronic system. With this tool, the system is built as a set of blocks, where each block can be described using different languages (C, VHDL, Verilog, System Verilog, OpenVera, etc.).

C) Implementation of the coder in C. Once the design of the structure of the test bench has been completed, the different blocks are being implemented in C. This modularity has allowed us to do an analysis about the most suitable coder for the image compression system. In the analysis a key factor is the compression rate.

D) Use and verification of the implemented wavelet transform. In the test bench, the block of the wavelet transform made in C has been replaced by its physical implementation (obtained from a synthesizable description VHDL).

E) Implementation of the system with a microprocessor. We have worked on the implementation of the coder block, using a microprocessor designed by us.

2.1.14 Objective 6.1.

A) Extension of automatic FPGA programming tools from high-level specifications. We have extended and improved our automatic FPGA programming tool based on high-level specifications (based on an object oriented language like SmallTalk-80) for SIMD applications. This tool has been tested over current FPGA based platforms and compared with commercial high-level synthesis tools. We have obtained better area results and comparable response time implementations, under the same experimental conditions, in the design of SIMD algorithms. Recently, we have improved the type inference technique used in our tool with backward data propagation, improving previous results. Also, we have started the formalization of our methodology in order to promote the use of these techniques in future commercial hardware compilers.

B) Extension and improvement of a Multimedia Router (MMR) prototype. We have extended the design of a Multimedia Router prototype, named as Simple MMR (SMMR), on an FPGA. We have evaluated its implementation in terms of required silicon area and maximum achievable frequency. Also, we have introduced some techniques to simplify some elements of SMMR like the Link Scheduler (LS) and the Switch Scheduler (SS). As a consequence of this work we have designed a new LS based on the temperature (thermometer) encoding of priority connections. This new LS maintains QoS (Quality of Service) guarantees of MMR router while reducing its implementation area and power consumption.

C) Proposal of a multimedia server with QoS based on FPGA. We have proposed a new link scheduler for a server with QoS for LAN/SAN and high performance cluster networks. The architecture of this link scheduler uses an FPGA to maintain and update the state of each connection. We have used credit-based flow control to avoid loosing data at the flit level as used in the MMR router. The first experimental results on a single Xilinx Virtex 2000 show that this architecture grows linearly with the number of connections but logarithmically with respect to the response time. We have tested configurations supporting from 4 to 128 connections in a Celoxica RC10000 platform.

D) Study of new FPGA organizations and high-speed internal lines for the design of NoC. We have carried out a first study of new FPGA organizations to evaluate their routing possibilities in the design of efficient NoC (Network on Chip). There are several possibilities to interconnect heterogeneous elements of current FPGA architectures. Their study will provide more efficient NoC designs, and help us to select the best topologies for a multiprocessor on chip. We have obtained first experimental results regarding the parameterization of routing resources for an FPGA in terms of latency and bandwidth.

2.2 Difficulties during project development

Objective 1.3 “Improving the performance of Internet routers” and Objective 2.1.A “Switch-fabric design for QoS-aware IP routers” have been canceled. These objectives were included due to the collaboration with the IBM Zurich Research Laboratory. Unfortunately, the crisis in the telecommunications market led IBM to close its Communications Department in Zurich, suddenly leading to the end of our collaboration. Thus, we decided to cancel Objective 1.3, but keeping its technical goals (enhanced switch architecture, congestion management), and developing them in a more generic form (thus making those goals to become part of Objective 1.1) or applying them to SANs (thus making those goals to become part of Objective 1.2). Regarding Objective 2.1A, we have decided to pursue a new research line aiming at reducing the complexity of cluster switches with QoS support.

Objective 2.4 "Analysis and dynamic parameter tuning for the medium access control layer in the IEEE 802.11 a/b WLAN". Due to the development of various WLAN technologies, we have undertaken the study of other WLAN technologies: HIPERLAN/2, IEEE 802.16, IEEE 802.11 g. As part of these efforts, a novel QoS scheme for HIPERLAN/2 has been proposed. Our results show the efficiency of the proposed scheme when supporting MPEG-4 video communications. We have introduced a framework comprising various resource request mechanisms to meet the quality of service requirements of multiple applications; and resource assignment mechanisms. Our numerical results show the effectiveness of the proposed framework when operating in a multiservice network.

Objective 4.1 “Integral task scheduling” has been postponed. Due to an extremely critical personal situation, the researcher who was implementing the simulator required for this objective had to temporarily leave his research tasks since July 2004. Due to the great effort required for another researcher to finish the implementation of such a complex tool, we decided to postpone this objective until the personal situation of the researcher allows him to return to his research tasks. We expect to accomplish some of the goals planned within this objective in a few months.

Objective 5.1. “Development of a quantizer and real-time video encoder amenable for implementation in a reconfigurable VLSI architecture”. Our hardware implementation uses integer numbers, so the quality of the image is not very high as a consequence of truncating the data.

Objective 6.1. Sub-objective entitled “Design of a parallel 3D wavelet on a cluster of PCs with FPGAs from high-level specifications” has been canceled. Other subgroups of this project have successfully advanced in the development of parallel and distributed algorithms, and hardware circuitry to compute the 3D wavelet. Moreover, there is an increasing interest in the evaluation of different elements to provide QoS on SAN and high performance clusters. Thus, we decided to cancel this sub-objective and to focus our efforts in the new sub-objective entitled “Proposal of a multimedia server with QoS based on FPGA”.

Overall, despite having canceled a few research goals, the current set of research objectives is more ambitious than the original one. Effectively, new research lines, not initially planned, have been

started, such as fault tolerance in InfiniBand, fault-tolerant routing for direct networks based on the use of intermediate nodes, a routing strategy based on interval routing, power consumption reduction in interconnection networks, and reduction of the number of necessary VCs for QoS in PC clusters.

3 Result indicators and main achievements

The project started one year and a half ago. Since then, we have filed an international joint patent with a large company in UK (Xyratex) and are close to file a second one. We have also filed an international joint patent titled "Perceptual cryptography in file-sharing environments". Also, 6 PhD students presented their dissertation, obtaining the maximum score. Moreover, we published a total of 138 papers in national and international conferences and journals. See Appendices for complete lists.

Additionally, we have become members of the European network of excellence named HiPEAC, have submitted two research proposals as workpackage leaders in the February 2005 call for Future and Emerging Technologies (FET) program on Advanced Computer Architecture (ACA), and have consolidated international collaborations and performed student and researcher exchanges with Simula Research Laboratory (Oslo, Norway), University of Southern California (Los Angeles, USA), Los Alamos National Laboratory (Los Alamos, USA), INRIA (Sophia Antipolis, France), CARV laboratory (Crete, Greece), Multimedia Communications Research Laboratory (University of Ottawa, Canada), Florida Atlantic University (Boca Raton, USA), and Virtual Reality Applications Center (Iowa State University, USA).

A Patents

Patent title: A Method of Congestion Management of a Network, a Signalling Protocol, A Switch, an End Station and a Network
Inventors: J. Duato, I. D. Johnson, F. Naven, J. Flich
Date: March 5, 2004
Country: USA

Patent title: Perceptual Cryptography in File-Sharing Environments

B PhD thesis

1. Aurelio Bermúdez, “Design of efficient subnet management mechanisms for InfiniBand networks”. PhD thesis. University of Castilla-La Mancha, September 2004.
2. Gregorio Bernabé García, “Diseño, Evaluación y Optimización de la Transformada Wavelet para Codificación de Vídeo Médico en Arquitecturas Monoprocesador”. PhD thesis. University of Murcia, November 2004.
3. Pedro Morillo Tena,“Mejora de las Prestaciones de los Entornos Virtuales Distribuidos”. PhD thesis. University of Valencia, December 2004.
4. Juan Piernas Cánovas, “Diseño, Implementación y Evaluación de un Sistema de Ficheros Transaccional con Separación de Datos y Metadatos”. PhD thesis. University of Murcia, October 2004.
5. Salvador Coll, “A strategy for efficient and scalable collective communication in the Quadrics network”. PhD thesis. Technical University of Valencia, July 2005.
6. Joaquin Cerdá, "Arquitectura VLSI de autómatas celulares para modelado físico". PhD thesis. Technical University of Valencia.

C Publication list

1. Manuel E. Acacio and José M. García. “Techniques for Improving the Performance and Scalability of Directory-based Shared-Memory Multiprocessors: A Survey”. Journal of Computer Science & Technology, Vol. 3, No. 2, pp. 1-8, Octubre 2003 (Invited paper).
2. M. E. Acacio, J. González, J. M. García and J. Duato. “A Two-level Directory Architecture for Highly Scalable cc-NUMA Multiprocessors”, IEEE Transactions on Parallel and Distributed Systems, Vol. 16, No. 1, pp. 67-79, January 2005.
3. M. E. Acacio, J. González, J. M. García and J. Duato. “An Architecture for High-Performance Scalable Shared-Memory Multiprocessors Exploiting On-chip Integration”. IEEE Transactions on Parallel and Distributed Systems, Vol. 15, No. 8, pp. 755-768, Agosto 2004.
4. F.J. Alfaro, J.L. Sánchez, Studying the Influence of the InfiniBand Packet Size to Guarantee QoS, The 10th IEEE Symposium on Computers and Communications (ISCC 2005). La Manga (SPAIN), 27-30 June 2005.

5. F.J. Alfaro, J.L. Sánchez, Tuning buffer size in InfiniBand to guarantee QoS, In proceedings of the 10th International Euro-Par Conference (EuroPar'04), Lecture Notes in Computer Science 3149, Springer, pp. 873-881, ISBN 3-540-22924-8, 2004

6. F. J. Alfaro, J. L. Sánchez, J. Duato, QoS in InfiniBand subnetworks, IEEE Transactions on Parallel and Distributed Systems, Vol. 15, n. 9, pp. 810-823, 2004.

7. M. Alonso, J. M. Martínez, V. Santonja, P. López, Reducing Power Consumption in Interconnection Networks by Dynamically Adjusting Link Width, Lecture Notes in Computer Science, Número: 3149 Número: 3149, Año: 2004, Páginas: 882-890 I.S.S.N.: 0302-9743

8. M. Alonso, J. M. Martínez, V. Santonja, P. López, J. Duato, Power Saving in Regular Interconnection Networks Built with High-Degree Switches, 19th International Parallel and Distributed Processing Symposium, IEEE Computer Society, Año: 2005 I.S.B.N.: 0-7695-2312-9

9. M. Alonso, J. M. Martínez, V. Santonja, P. López, J. Duato, Reducción del consumo de potencia en redes de interconexión construidas con conmutadores de alto grado, Enviado a: XVI Jornadas de Paralelismo, Año: 2005.

10. J.L. Aragón, D. Nicolaescu, A. Veidenbaum and A. M. Badulescu. "Energy-Efficient Design for Highly Associative Instruction Caches in Next-Generation Embedded Processors". In Proceedings of the International Conference on Design, Automation and Test in Europe (DATE), Paris, France, Feb. 2004.

11. J.L. Aragon and A. Veidenbaum. "Low-power Fetch Unit Design for Superscalar Processors". In Proceedings of the XV Jornadas de Paralelismo, Almeria (Spain), September 2004.

12. V. Arnau, R. Avendaño, Simulador de Técnicas de Predicción Dinámicas", XV Jornadas de Paralelismo, pp. 460-464. I.S.B.N.: 84-8240-714-7. Septiembre de 2004.

13. V. Arnau, I. Marín. Fast analysis of highly repetitive sequences in human chromosomes using a novel search algorithm, WSEAS Transactions on Biology and Biomedicine, pp. 107-111. ISSN: 1109-9518. Grecia, Enero de 2004.

14. A. Bermúdez, R. Casado, F. J. Quiles, J. Duato, Use of Provisional Routes to Speed-up Change Assimilation in InfiniBand Networks, IEEE International Workshop on Communication Architecture for Clusters (CAC'04), Santa Fe (NM), USA, IEEE Computer Society, ISSN 1530-2075, ISBN 0-7695-2132-0, pp. 186-196. Abril 2004.

15. A. Bermúdez, R. Casado, F. J. Quiles, J. Duato, Distributing InfiniBand Forwarding Tables, 10th International EuroPar Conference, Pisa, Italia, Springer-Verlag, ISSN 0302-9743, ISBN 3-540-22924-8, pp. 864-872, Septiembre 2004.

16. A. Bermúdez, R. Casado, F. J. Quiles, J. Duato, Fast Routing Computation on InfiniBand Networks. Aceptada su publicación en la revista IEEE Transactions on Parallel and Distributed Systems, IEEE Computer Society Press, 2005.

17. G. Bernabé, J. M. García y J. González. "An Efficient 3D Wavelet Transform on Hyper-Threading Technology". Technical Report. UM-DITEC-2004-02.

18. G. Bernabé, J. M. García y J. González. “Reducing 3D Fast Wavelet Transform Execution Time Using Blocking and the Streaming SIMD Extensions”. Journal of VLSI Signal Processing Systems, Volume 41, Number 2, September 2005.

19. A. Bierbaum, P. Hartling, P. Morillo and C. Cruz-Neira, Immersive Clustering with VR Juggler, in Int. Conf. in Computational Science and Its Applications (ICCSA’ 2005), Singapore, May 2005. Springer Verlag.

20. B. Caminero, C. Carrión, F. J. Quiles, J. Duato y S. Yalamanchili, Traffic Scheduling Solutions with QoS Support for an Input-Buffered MultiMedia Router. Aceptada su publicación en la revista IEEE Transactions on Parallel and Distributed Systems, 2005.

21. Óscar Cánovas, Gabriel López, Antonio F. Gómez-Skarmeta. “A Credential Conversion Service for SAML-based scenarios”. Proceedings of First European PKI Workshop. Lecture Notes in Computer Science 3093, Springer Verlag, pp. 297-305, Samos (Greece), Junio 2004.

22. M. Canseco, J. M. Claver, G. León, I. Vilata. “Primeras experiencias en la implementación de un encaminador QoS sobre una FPGA”. FPGAS: Computación & Aplicaciones. JCRA’04, pp. 289-296 . ISBN 84-688-7667-4. Almería, September, 2004.

23. M. Canseco, J. M. Claver, G. León, I. Vilata. “Prototipado de un MMR Simple en una FPGA”. XV Jornadas de Paralelismo ’04, pp. 345-350. ISBN 84-8240-714-7. Almería, September, 2004.

24. R. Casado. Disponibilidad de la red, Curso de Verano UCLM: Una ventana abierta hacia el futuro de la arquitectura de computadores, Universidad de Castilla-La Mancha , ISBN 84-921873-6-0, pp. 95-120. July 2004.

25. J. Cerdá, R. Gadea, J.D. Martinez and A. Sebastiá. "A tool for implementing and exploring SBM models: universal 1D invertible cellular automata". Lecture Notes in Computer Science: Mechanisms, Symbols, and Models Underlying Cognition, No. 3561, pp. 279-289, 2005.

26. J. Cerdá, O. Amoraga, R. Torres, R. Gadea and A. Sebastiá. "FPGA implementations of the RNR cellular automata to model electrostatic field". Lecture Notes in Computer Science: High Performance Computing for Computational Science, No. 3402, pp. 382-395, 2004.

27. J. Cerdá, R. Gadea and G. Paya. "Implementing a Margolus neighborhood cellular automata on a FPGA". Lecture Notes in Computer Science: Artificial Neural Nets Problem Solving Methosds, No. 2687, pp. 121-128, 2003.

28. J. Cerdá, R. Gadea, V. Herrero and A. Sebastiá. "On the Implemention of a Margolus neighborhood cellular automata on FPGA". Lecture Notes in Computer Science: Field-Programmable Logic and Applications, No. 2778, pp. 776-785, 2003.

29. J. M. Claver, M.C. Carrión, M.Canseco, M.B. Caminero, and F.J. Quiles. “A New Efficient Link Sheduling Algorithm to Guarantee QoS on Clusters”. Euro-Par 2005. Lisboa, Portugal, 8/30-9/2, 2005. Accepted and pending to be published in Lecture Notes in Computer Science (LNCS).

30. S. Coll, J. Duato, F. J. Mora, F. Petrini, and A. Hoisie, Collective Communication Patterns on the Quadrics Interconnection Network, in V. Getov, M. Gerndt, A. Hoisie, A. Malony, and B. Miller (Ed.), Performance Analysis and Grid Computing, Kluwer, pp. 93-107, ISBN 1-4020-7693-2, 2004.

31. R.J. Colom, R. Gadea, F.J. Ballester and M. Martinez. "Flexible architecture for the implementation of the two-dimensional discrete wavelet transform (2D-DWT) oriented to FPGA devices". Microprocessors and Microsystems, No. 18, pp. 509-518, 2004.

32. P. Cuenca, L. Orozco-Barbosa, F. Delicado, A. Garrido, Breakpoint Tuning in DCT-Based Nonlinear Layered Video Codecs, Vol. 16.pp. 2555-2570. Journal on Applied Signal Processing. November 2004.

33. F. Delicado, P. Cuenca, L. Orozco-Barbosa, Developing a QoS Framework for Media Streaming over TDMA/TDD Wireless Networks, Journal of Wireless and Mobile Computing. 2005.

34. F. Delicado, P. Cuenca, L. Orozco-Barbosa, A. Garrido, Design and Evaluation of a QoS-aware Framework for Wireless TDMA-TDD Networks, Wireless Personal Communications Journal. 2005

35. F. Delicado, P. Cuenca, L. Orozco-Barbosa, Multiservice Communications over TDMA/TDD Wireless LANs, Lecture Notes in Computer Science. Vol. 3510, pp. 107-116, 2005.

36. F. Delicado, P. Cuenca, L. Orozco-Barbosa, QoS-aware Resource Request Mechanisms for HIPERLAN/2, IEEE Global Telecommunications Conference (GLOBECOM 2004). December, 2004

37. F. Delicado, P. Cuenca, L. Orozco-Barbosa, A QoS-aware WLAN Resource Request Mechanism for Delay Sensitive Traffic, IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC-04). September, 2004

38. F. Delicado, P. Cuenca, L. Orozco-Barbosa, A QoS-aware Resource Request Mechanism for Delay Sensitive Services over TDMA/TDD Wireless Networks, 1st International Conference on E-business and Telecommunication Networks. August, 2004

39. F. Delicado, P. Cuenca, L. Orozco-Barbosa and A. Garrido, A Class-based Allocation Mechanism for Delay Sensitive Traffic in WLAN, 18 th International Conference on Advanced Information Networking and Applications. Fukuoaka, Japan. March 2004.

40. J. Duato, J. Flich, T. Nachiondo, A Cost - Effective Technique to Reduce HOL Blocking in Single- Stage and Multistage Switch Fabrics , presentado en 12th Euromicro Conference on Parallel, Distributed and Network - based Processing, Febrero, 2004, A Coruña.

41. J. Duato, I. Johnson, J. Flich, F. Naven, P. J. García, T. Nachiondo, A New Scalable and Cost - Effective Congestion Management Strategy for Lossless Multistage Interconnection Networks, presentado en 11th International Symposium on High- Performance Computer Architecture, San Francisco, EEUU, Febrero 2005.

42. J. Duato, O. Lysne, R. Pang and T. M. Pinkston, A Theory for Deadlock-Free Dynamic Network Reconfiguration, IEEE Transactions on Parallel and Distributed Systems 16(5), 2005.

43. G. Fernández-Escribano, P. Cuenca, L. Orozco-Barbosa, A. Garrido, A Fast Intra-Frame prediction algorithm for MPEG- 2/H.264 video transcoders, Proceedings of the IEEE International Conference on Image Processing, ICIP-05, September 11-14, 2005, Genova, Italy.

44. G. Fernández-Escribano, P. Cuenca, L. Orozco-Barbosa, A. Garrido, Computational complexity reduction of Intra-Frame prediction in MPEG-2/H.264 video transcoders Proceedings of the IEEE International Conference on Multimedia & Expo, ICME-05, July 6-8 de 2005, Amsterdam, Holanda.

45. G. Fernández-Escribano, J. Villalón, P. Cuenca, L. Orozco-Barbosa, A. Garrido, On the Capabilities of Intra-Frame prediction in H.264 video encoders Proceedings of the IEEE Pacific Rim Conference on Communications, Computers and Signal Processing, PacRIM-05, Victoria, B.C., Canada, August 24-26 de 2005.

46. G. Fernández-Escribano, P. Cuenca, L. Orozco-Barbosa, A. Garrido, FIMDA: A Fast Intra-Frame Mode Decision Algorithm for MPEG-2/H.264 Transcoding Lecture Notes in Computer Science (LCNS), Advanced Concepts for Intelligent Vision Systems, ACIVSï05 (in press), September 20-23, 2005, Antwerp, Belgium.

47. G. Fernández-Escribano, P. Cuenca, L.Orozco-Barbosa and A. Garrido, Transcodificación Heterogénea de Video MPEG-2/H.264: Cuestiones Abiertas, Jornadas de Paralelismo. Actas de las Jornadas de Paralelismo.pp. 132-137. September 2004

48. R. Fernández, G. Bernabé, J. M. García y M. E. Acacio. "Codificador de vídeo basado en Wavelet 3D usando OpenMP y Pthreads". Actas de las XV Jornadas de Paralelismo, Universidad de Almería, ISBN: 84-8240-714-7, Almería-España, Septiembre 2004.

49. R. Fernández, G. Bernabé, J. M. García y M. E. Acacio. "Optimizing a 3D-FWT Video Encoder for SMPs and HyperThreading Architectures". 13th EUROMICRO Conference on Parallel, Distributed and Network-Based Processing, 2005.

50. Ricardo Fernández y José M. García. RSIMx86: A Cost Effective Performance Simulator. The High Performance Computing & Simulation (HPC&S) Conference. Riga (Latvia), Junio 2005.

51. Juan Fernández, Eitan Frachtenberg, Fabrizio Petrini, Kei Davis and José Carlos Sancho. "Architectural Support for System Software on Large-Scale Clusters". Actas de International Conference on Parallel Processing (ICPP´04), Montreal, Québec (Canadá). Agosto de 2004.

52. Juan Fernández, Eitan Frachtenberg, Fabrizio Petrini, Kei Davis and José Carlos Sancho. "On the Design of System Software for Large-Scale Clusters". Actas de XV Jornadas de Paralelismo. Almería. Septiembre de 2004.

53. Juan Fernández, Fabrizio Petrini and Eitan Frachtenberg. "Monitoring and Debugging Parallel Software with BCS-MPI on Large-Scale Clusters". Workshop on System Management Tools for Large-Scale Parallel Systems. International Parallel and Distributed Processing Symposium, Denver, CO (USA). Abril de 2005.

54. Juan Fernández, Fabrizio Petrini and Eitan Frachtenberg. "Achieving Predictable Performance with BCS-MPI". Capítulo del libro "Engineering the Grid". Jack Dongarra, Hans Zima, Adolfo Hoisie, Laurence Yang and Beniamino di Martino editors. Pendiente de publicación en 2005.

55. J. L. Ferrer, E. Baydal, A. Robles, P. López, and J. Duato, MVCC: An Efective Marking/Validation-based Control Management Strategy for InfiniBand Networks, XVI Jornadas de Paralelismo, Septiembre 2005.

56. Eitan Frachtenberg, Kei Davis, Fabrizio Petrini, Juan Fernández y José Carlos Sancho. “Designing Parallel Operating Systems via Parallel Programming”. Actas de Euro-Par 2004, LNCS 3149, pp. 689-696. Ed. Springer-Verlag. Pisa, Italia. Septiembre de 2004.

57. Eitan Frachtenberg, Fabrizio Petrini, Dror G. Feitelson and Juan Fernández. “Adaptive Parallel Job Scheduling with Flexible CoScheduling”. IEEE Transactions on Parallel and Distributed Systems. Pendiente de publicación en 2005.

58. R. Gadea, R.J. Colom, J. Cerda and A. Sebastia. "FPGA implementation of a pipelined on-line backpropagation". Journal of VLSI signal processing, No. 40, pp. 189-213, 2005.

59. R. Gadea, A. Ramírez, J. Cerdá and R. Colom. "FPGA Implementation of Adaptive Non-Linear Predictors for Video Compression". Lecture Notes in Computer Science: Field-Programmable Logic and Applications, No. 2778, pp. 1016-1019, 2003.

60. Pablo E. García, Juan Fernández, Fabrizio Petrini and José M. García. “Cluster Computing with QsNetII”. Actas de XVI Jornadas de Paralelismo. Congreso Español De Informática (CEDI´ 2005). Granada. Septiembre de 2005.

61. Félix J. García, Gregorio Martínez, Óscar Cánovas, Antonio F. Gómez-Skarmeta. “A proposal of a CIM-based Policy Management Model for the OGSA Security Architecture”. In Proceedings of GADA 2004, Lecture Notes in Computer Science 3292, Springer, pp. 165-174, Agia Napa (Cyprus), October 2004.

62. F. J. García, O. Cánovas, G. Martínez, y A.F. Gómez-Skarmeta: “Self-configuration of grid nodes using a policy-based management architecture”, APGAC Workshop, ICCS, 6/2004.

63. I. García, R. Mollá and P.Morillo, From Continous to Discrete Games, Proc. of 2004 Computer Graphics International (CGI’04), pp. 626-630, IEEE Computer Society Press, Crete, Greece. June, 2004.

64. P. J. García, J. Flich, J. Duato, F.J. Quiles, I. Johnson, F. Naven, On the Correct Sizing on Meshes through an Effective Congestion Management Strategy , aceptado para presentar en Euro- Par 2005, Lisboa, Agosto 2005.

65. G. Garcia-Mateos, A. Garcia-Meroño, C. Vicente, A. Ruiz, Pedro E. Lopez-de-Teruel, “Time and Date OCR in CCTV Video”, Proceedings of the 13th International Conference on Image Analysis and Processing (ICIAP 2005), Cagliari (Italy), 2005.

66. M. Giménez, J. M. Benlloch, J. Cerdá, B. Escat, M. Fernández, E. N. Giménez, C. W. Lerche, J.D. Martínez, F.J. Mora, N. Pavón, F. Sánchez and A. Sebastià “Medium field of view multiflat panel-based portable gamma camera”, Nuclear Instruments and Methods. ISSN: 0168-9002. Vol: 525 Issues I-2 pags. 298-302, 2004.

67. M. E. Gómez, J. Duato, J. Flich, P. Lopez, A. Robles, N.A. Nordbotten, T. Skeie, O. Lysne. “A New Adaptive Fault-Tolerant Routing Methodology for Direct Networks,” International Conference on High Performance Computing (HiPC), 2004.

68. M. E. Gómez, J. Flich, P. López, A. Robles, J. Duato, N.A. Nordbotten, O. Lysne, T. Skeie. “An Effective Fault-Tolerant Routing Methodology for Direct Networks,” in Proc. International Conference on Parallel Processing (ICPP), 2004, pp. 222-231.

69. M. E. Gómez, J. Duato, J. Flich, P. López, and A. Robles, N.A. Nordbotten, O.Lysne, T. Skeie. “An Efficient Fault-Tolerant Routing Methodology for Meshes and Tori,” Computer Architecture Letters, Vol. 3, May 2004.

70. M. E. Gómez, P. Lopez, J. Duato. “ A Memory-Effective Routing Strategy for regular Interconnection Networks”. Proc. International Parallel and Distributed Processing. Best paper award. Denver (USA), 2005.

71. M. E. Gómez, P. Lopez, J. Duato. “ A Memory-Effective Fault-Tolerant Routing Strategy for Direct nterconnection Networks”. International Symposium on Parallel and Distributed Computing Proc. International Parallel and Distributed Processing. Lille (Francia), 2005.

72. F. Grimaldo, M. Lozano, F. Barber, J. M. Orduña, Integrating Social Skills in Task-Oriented 3-D IVA, to appear in Proceedings of 5TH International Working Conference on Intelligent Virtual Agents, 2005.

73. M. Koibuchi, J. C. Martinez, J. Flich, A. Robles, P. López, and J. Duato, Enforcing In-Order Packet Delivery in PC Clusters using Adaptive Routing Information, 2004 Workshop on High Performance Computing and Networking (HPCN-04), Tokyo, Japón, 2004

74. M. Koibuchi, J. C. Martinez, J. Flich, A. Robles, P. López, and J. Duato, “Enforcing In-Order Packet Delivery in System Area Networks with Adaptive Routing”, To appear in Special Issue on Design and Performance of Networks for Super, Cluster and Grid Computing to appear in the Journal of Parallel and Distributed Computing (JPDC), 2005.

75. G. León, J. M. Claver, G. Fabregat. “Optimizing Area on the Generation of Specific Circuits in FPGAs for SIMD Applications”. International Workshop on Applied Reconfigurable Computing 2005, pp. 160-167. ISBN 972-99353-8-6. Algarve, Portugal, February 22-23, 2005.

76. G. León, J. M. Claver, G. Fabregat. “Optimizing Area on the Automatic Generation of Circuits in FPGAs for SIMD Applications”. XV Jornadas de Paralelismo ’04, pp. 438-443. ISBN 84-8240-714-7. Almería, September, 2004.

77. Ch. W. Lerche, J. M. Benlloch, F. Sánchez, N. Pavón, N. Giménez, M. Fernández, M. Giménez, A. Sebastiá, J. Martínez and F.J. Mora. “Depth of interaction detection with enhanced position-sensitive proportional resistor network”, Nuclear instruments and methods in physics research ISSN: 0168-9002. Vol.: 537 Issue 1-2. Pags. 326-330, 2005

78. Gabriel López, Óscar Cánovas, Antonio F. Gómez-Skarmeta. “Servicio de acceso a la red basado en autorización SAML”. En Avances en Criptología y Seguridad de la Información, Actas de la VIII Reunión Española sobre Criptología y Seguridad de la Información, pp. 625-635, Leganés (España), Septiembre 2004.

79. Gabriel López, Oscar Cánovas, Antonio F. Gómez Skarmeta, Rafael Marín. "A Network Access Control Approach based on the AAA Architecture and Authorization Attributes”, International Workshop on System and Network Security. In conjunction with IEEE 19th IPDPS.

80. Gabriel López; Cánovas, O., Gómez, A. F., y Marín, R. “Different Network Access Control Methods based on the AAA Architecture.” Journal of Network and Computer Applications. Aceptado. Pendiente de publicación. 2005.

81. Gabriel López; Cánovas,O., Gómez, A. F., Otenko, O., y Chadwick, D.W. “A heterogeneous network access service based on PERMIS and SAML.” Proceedings of 2nd European PKI Workshop, Canterbury (UK). Aceptado. 2005.

82. O. Lysne, T. M. Pinkston and J. Duato, A Methodology for Developing Deadlock-Free Dynamic Network Reconfiguration Processes, IEEE Transactions on Parallel and Distibuted Systems 16(5), 2005.

83. A. Martínez, F. J. Alfaro, J. L. Sánchez, J. Duato, A strategy to reduce the number of necessary VCs for QoS support in clusters. Accepted for HiPC 2005.

84. J. C. Martínez, J. Flich, A. Robles, P. López, and J. Duato, Supporting Adaptive Routing in InfiniBand Networks, Journal of System Architecture. No. 49 (2003), Elsevier, pp. 441-456. ISSN: 1383-762/0165-6074.

85. J. C. Martínez, J. Flich, A. Robles, P. López, J. Duato, Solving out-of-order packet delivery in interconnection networks using adaptive routing, XV Jornadas de Paralelismo, Almería 2004.

86. J. C. Martínez, M. Koibuchi, J. Flich, A. Robles, P. López, and J. Duato, In-Order Packet Delivery in Interconnection Networks using Adaptive Routing, Aceptado en 19th International Parallel & Distributed Processing Symposium (IPDPS’05), Denver (EEUU), 2005.

87. J.D. Martínez, A. Sebastiá, J. Cerdá, R. Esteve, F.J. Mora, J.F. Toledo, J.M. Benlloch, N. Giménez, M. Giménez, Ch. W. Lerche, “Data acquisition electronics for positron emission mammography (PEM) detectors”, Nuclear Instruments and Methods In Physics Research ISSN: 0168-9002 . Vol: 537 Issue 1-2. Pags. 335-338, 2005.

88. M. Martínez, F.J. Ballester, G. Paya, R. Colom, R. Gadea and J. Belenguer. "FPGA Custom DSP for ECG Signal Analysis and Compression". Lecture Notes in Computer Science: Field-Programmable Logic and Applications, No. 3203, pp. 954-958, 2004.

89. R. Martínez, J.L. Sánchez, F.J. Alfaro, V. Chirivella, J. Flich, Studying the effect of the design parameters on the interconnection network performance in NOWs, In proceedings of the 13th Euromicro Conference on Parallel, Distributed and Network-based Processing (PDP’98), IEEE Computer Society Press, pp. 102-109, ISBN 0-7695-2280-7, 2005.

90. F. Micó, P. Cuenca, L. Orozco-Barbosa, QoS in IEEE 802.11 wireless LAN: current research activities, in Proc. of IEEE Canadian Conference on Electrical and Computer Engineering, pp. 447- 452. May 2004.

91. F. Micó, P. Cuenca, and L. Orozco-Barbosa, QoS Mechanisms for IEEE 802.11 Wireless LANs, Lecture Notes in Computer Science, Vol. 3079, pp. 609-623. ISSN: 0302-9743. 2004.

92. J. M. Montañana, J. Flich, A. Robles, P. López, and J. Duato, A Transition-Based Fault-Tolerant Routing Strategy for InfiniBand Networks, Proceedings of the Workshop on Communication Architecture for Clusters (CAC’04), April 2004.

93. J. M. Montañana, J. Flich, A. Robles, P. López, and J. Duato, Providing Fault Tolerance To InfiniBand Networks, Actas de las XV Jornadas de Paralelismo, Septiembre 2004.

94. J. M. Montañana, J. Flich, A. Robles, and J. Duato, An Scalable Methodology for Computing Fault-free Paths in InfiniBand Torus Networks, Submitted to 6th International Symposium on High Performance Computing (ISHPC-VI), Japan, September 2005.

95. J. M. Montañana, J. Flich, A. Robles, and J. Duato, Reachability-Based Fault-Tolerant Routing Methodology, XVI Jornadas de Paralelismo, Septiembre 2005.

96. P. Moreno, P. Morillo, J. M. Orduña, and J. Duato, Calidad de Servico en Entornos Virtuales Distribuidos, accepted for presentation in XVI Jornadas de Paralelismo, I Congreso Español de Informática. Granada (SPAIN), 2005.

97. P. Morillo, J. M. Orduña, M. Fernández, and J. Duato, A Comparison Study of Metaheuristic Techniques for Providing QoS to Avatars in DVE Systems, in Proceedings of the International Conference on Computational Science and its Applications 2004 (ICCSA'04), pp. 661-670, ISBN 3-540-22056-9, May 2004.

98. P. Morillo, J.M. Orduña, M. Fernández, J. Duato, A New Partitioning Approach for Distributed Virtual Environment Systems, Proceedings of XV Jornadas de Paralelismo, pp. 12-17. ISBN 84-8240-714-7. 2004.

99. P. Morillo, J.M. Orduña, M. Fernández, J. Duato, Metaheuristic Solutions for Providing QoS in DVE Systems, Proceedings of XV Jornadas de Paralelismo, pp. 18-23. ISBN 84-8240-714-7. 2004.

100. P. Morillo, J. M. Orduña, M. Fernández, and J. Duato, A Fine-Grain Method for Solving the Partitioning Problem in Distributed Virtual Environment Systems, in Proceedings of the IASTED International Conference on Parallel and Distributed Computing and Systems (PDCS 2004), Iasted Acta Press, pp. 292-297, ISBN 0-88986-421-7, 2004.

101. P. Morillo, J.M. Orduña and M. Fernández, A Comparison Study of evolutive algorithms for Solving the Partitioning Problem in Distributed Virtual Environment Systems, Parallel Computing, Vol. 30, No. 5, pp. 585-610, 2004.

102. P. Morillo, P. López, J.M. Orduña and M. Fernández, A New Genetic Approach for the Partitioning Problem in Distributed Virtual Environment Systems, Lecture Notes in Artificial Intelligence, Vol. 3040, pp. 76-85, 2004.

103. P. Morillo, J. M. Orduña, M. Fernández, and J. Duato, A Method for Providing QoS in Distributed Virtual Environment Systems, in Proceedings of the EuroMicro Conference on Parallel, Distributed and Network-Based Processing (PDP 2005), IEEE Computer Society Press, pp. 152-159, ISBN 0-7695-2280-7, 2005.

104. P. Morillo, J. M. Orduña, M. Fernández, and J. Duato, Improving the Performance of Distributed Virtual Environment Systems, IEEE Transactions on Parallel and Distributed Systems, Vol., No 7, pp. 637-649, 2005.

105. P. Morillo, M. Fernández, and J. M. Orduña, Solving the Partitioning Problem in Distributed Virtual Environment Systems Using Evolutive Algorithms, Handbook of Bioinspired Algorithms and Applications, chapter 30, pp. 529-552, Chapman&Hall, 2005.

106. T. Nachiondo, J. Flich, J. Duato, Efficient Reduction of HOL Blocking in Multistage Networks , presentado en 2005 Workshop on Communication Architecture for Clusters, dentro del 19th International Parallel and Distributed Processing Symposium, Abril 2005, Denver, USA.

107. T. Nachiondo, J. Flich, J. Duato, M. Gussat, Cost /Performance Trade - offs and Fairness Evaluation of Queue Mapping Policies, aceptado para presentar en Euro- Par 2005, Lisboa, Agosto 2005.

108. N. A. Nordbotten, M. E. Gómez, J. Flich, P. López, A. Robles, T. Skeie, O. Lysne, J. Duato. "A Fully Adaptive Fault-Tolerant Routing Methodology Based on Intermediate Nodes," IFIP International Conference on Network and Parallel Computing (NPC), 2004.

109. J. M. Orduña, F. Silla, and J. Duato, On the Development of a Communication-Aware Task Mapping Technique, Journal of Systems Architecture, Vol. 50, No. 4, pp. 207-220, 2004.

110. Alejandro Orozco, Manuel E. Acacio and José M. García. "Estudio y Evaluación del Encaminamiento Multidestino en Arquitecturas cc-NUMA con Directorios Comprimidos". Actas de las XVI Jornadas de Paralelismo, Septiembre 2005.

111. L. Orozco-Barbosa, M. Toukourou, P. Cuenca, A. Miri, A QoS Framework for Scalable MPEG-2 Video-on-Demand Services over IP-based Networks, Journal of Internet Technology. Vol. 5(4). pp.389-403. October 2004.

112. G. Paya, M. Martinez, F.J. Ballester and F.J. Mora. "Fully parameterized discrete wavelet packet transform architecture oriented to FPGA". Lecture Notes in Computer Science: Field-Programmable Logic and Applications, No. 2778, pp. 533-542, 2003.

113. Fabrizio Petrini, Adam Moody, Juan Fernández, Eitan Frachtenberg and Dhabaleswar K. Panda. "NIC-based Reduction Algorithms for Large-Scale Clusters". International Journal of High Performance Computing and Networking (IJHPCN). Pendiente de publicación en 2005.

114. J. Piernas, T. Cortes, and J. M. García. "Traditional File Systems versus DualFS: a Performance Comparison Approach". IEICE Transactions on Information and Systems, Vol.E87-D No.7 p. 1703-1711, July 2004.

115. R. Risueño, P. Cuenca, F. Delicado and L.Orozco-Barbosa, On the Effect of Handover Mechanisms on the Performance of Video Communications in WATM Networks, Journal of Wireless and Mobile Computing. 2005 (in press).

116. R. Risueño, P. Cuenca, F. Delicado, L. Orozco-Barbosa, and A. Garrido, On the Traffic Disruption Time and Packet Lost Rate during the Handover Mechanisms in Wireless Networks, 18 th International Conference on Advanced Information Networking and Applications. Fukuoaka, Japan. March 2004.

117. R. Risueño, P. Cuenca, F. Delicado, L. Orozco-Barbosa, and A. Garrido, On the Effect of the Handover Mechanisms in QoS Performance in Wireless Multimedia Networks, 6th International Workshop on Multimedia Network Systems and Applications. Tokyo, Japan. March 2004.

118. Alberto Ros, Manuel E. Acacio and José M. García. "A Novel Lightweight Directory Architecture for Scalable Shared-Memory Multiprocessors". Proc. of the Euro-Par 2005, LNCS, Agosto 2005.

119. Alberto Ros, Manuel E. Acacio and José M. García. “Diseño y Evaluación de una Arquitectura de Directorio Ligero para Multiprocesadores de Memoria Compartida Escalables”. Actas de las XVI Jornadas de Paralelismo, Septiembre 2005.

120. S. Rueda, P. Morillo, J. M. Orduña, and J. Duato, A Sexual Elitist Genetic Algorithm for Providing QoS in Distributed Virtual Environment Systems, Proc. of 2005 Int. Parallel and Distributed Processing Symposium (IPDPS’ 2005). ISBN: 0-7695-2312-9. 2005.

121. S. Rueda, P. Morillo, J. M. Orduña, and J. Duato, Algoritmos Genéticos para Proporcionar Calidad de Servicio en Entornos Virtuales Distribuidos, accepted for presentation in XVI Jornadas de Paralelismo, I Congreso Español de Informática. Granada (SPAIN), 2005.

122. F. Sanchez, J. M. Benlloch, B. Escat, N. Pavón, F.J. Mora, A. Sebastià, E. Porras, D. Kadi-Hanifi, J. A. Ruiz, “Design and tests of a portable mini gamma camera”, Medical Physics. ISSN: 0094-2405. Vol. 31 Nº 6. Pag. 1384-1397, 2004.

123. Manuel Sánchez, Óscar Cánovas, Diego Sevilla. “Firma de trabajos en la integración Globus 2 y Globus 3”. In Avances en Criptología y Seguridad de la Información, Actas de la VIII Reunión Española sobre Criptología y Seguridad de la Información, pp. 341-352, Leganés (Spain), September 2004.

124. Manuel Sánchez, Óscar Cánovas, Diego Sevilla, Antonio F. Gómez-Skarmeta. A Service-Based Architecture for Integrating Globus 2 and Globus 3. In Advances in Grid Computing - EGC 2005. LNCS 3470. The Netherlands, February 14-16, 2005.

125. José Carlos Sancho, Fabrizio Petrini, Greg Jonhson, Juan Fernández y Eitan Frachtenberg. “On the Feasibility of Incremental Checkpointing for Scientific Computing”. Actas de International Conference on Parallel and Distributed Processing Symposium (IPDPS´04), Santa Fe, NM (USA). Abril de 2004.

126. J. C. Sancho, A. Robles, J. Duato, An effective methodology to compute up*/down* routing tables in NOWs, IEEE Transaction on Parallel and Distributed Systems Vol. 15, No. 6, 2004, IEEE Computer Society, ISSN: 1045-9219.

127. T. Skeie, O. Lysne, J. Flich, P. López, A. Robles, J. Duato, LASH-TOR: A Generic Transition-Oriented Routing Algorithm, 10th International Conference on Parallel and Distributed Systems (ICPADS 2004), IEEE Computer Society Press, Newport Beach, CA, USA, July 2004, ISBN: 0-7695-2152-5.

128. Francisco J. Villa, Manuel E. Acacio and José M. García. “Evaluating IA-32 Web Servers through Simics: A Practical Experience”. Journal of Systems Architecture, Vol. 51 (2005) 251-264.

129. Francisco J. Villa, Manuel E. Acacio and José M. García. “Thread-Level Speculation and Chip-Multiprocessors: An Overview”. Proc. of the PARA´04 Workshop on State-of-the-Art Scientific Computing, Lyngby (Dinamarca), Junio 2004.

130. Francisco J. Villa, Manuel E. Acacio and José M. García. “On the Evaluation of x86 Web Servers Using Simics: Limitations and Trade-Offs”. Proc. of the 4th International Conference on Computational Science, LNCS 3036, pp. 541-544, Cracovia (Polonia), Junio 2004.

131. Francisco J. Villa, Manuel E. Acacio and José M. García. “Evaluación de un Servidor Web Multiprocesador mediante Simics”. Actas de las XIV Jornadas de Paralelismo, Septiembre 2003.

132. Francisco J. Villa, Manuel E. Acacio and José M. García. “Memory Subsystem Characterization in a 16-core Snoop-Based Chip-Multiprocessor Architecture”. 2005 International Conference on High Performance Computing and Communications. Septiembre, 2005.

133. Francisco J. Villa, Manuel E. Acacio and José M. García. “Performance Evaluation of a Bus-Based 16-core Chip-Multiprocessor”. Actas de las XVI Jornadas de Paralelismo, Septiembre 2005.

134. J. E. Villalobos, J. L. Sánchez, J. A. Gámez, J. C. Sancho, A. Robles, A Methodology to Evaluate the Effectiveness of Traffic Balancing Algorithms, Lecture Notes in Computer Science. No. 3149. Springer-Verlag, 2004, ISBN: 3-540-22924-8.

135. J. Villalón, P. Cuenca and L. Orozco-Barbosa, QoS Provisioning Mechanisms for IEEE 802.11 WLAN: A Performance Evaluation, Proc. of 10th IEEE/IFIP International Conference on Personal Wireless Communications PWC’05. Colmar, August 2005.

136. J. Villalón, F. Mico, P. Cuenca and L. Orozco-Barbosa, QoS Support for Time-Constrained Multimedia Communications in IEEE 802.11 WLANs: a Performance Evaluation, Proc. of IEEE International Conference on Multimedia Communications Systems (ICMCS-05). Montreal, August 2005.

137. J. Villalón, P. Cuenca and L. Orozco-Barbosa, Limitations and Capabilities of QoS Support in IEEE 802.11 WLANs, Proc. of IEEE Pacific Rim Conference on Communications, Computers and Signal Processing, PACRIM-05. Victoria, August 2005.

138. J. Villalón, P. Cuenca and L. Orozco-Barbosa, On the Effectiveness of IEEE 802.11e QoS Support in Wireless LAN: A Performance Analysis, Submitted to the International Conference on High Performance Computing and Communications, HPCC-05, Naples, Italy, September 2005.