

Conferencias

Conferencistas	Temas	Idioma	Session Chairs
Francisco Tirado	Multithreading, multi/manycores	Español	Tendencias actuales en el diseño de procesadores
Marc Massot	Multiphase reacting flows	Inglés	Modeling issues to improve performance (UNM)
Nelson Maculan	Investigación Operativa y	Español	Toma de Decisiones - Carlos Pozuelo (CICESE)
Mauricio Marin	Scalable Distributed Web	Español	Engines Juan de Dios Garbett (UAA)
Fabio Casati	Liquid Journals: how social	Inglés	computing and liquid knowledge will shape and navigate it
Luca Cernuzzi (UCA)			
Walter Bender	Software Libre and the	Inglés	Future of Learning
Benjamín Barán (UNA)			

Resúmenes

1. Multithreading, multi/manycores. Tendencias actuales en el diseño de procesadores.

Francisco Tirado

AULA MAGNA FP-UNA, MIÉRCOLES 20 DE OCTUBRE, 18:30 HS.

Desde su aparición hace casi 50 años, los procesadores han doblado su velocidad cada 18 meses. Esto ha sido debido, por una parte, a los avances en la integración de circuitos que permiten duplicar el número de transistores cada 18 meses y una mayor velocidad del reloj y por otra, a mejoras de su arquitectura (organización interna de los diferentes módulos operativos) orientadas a ejecutar un número cada vez mayor de instrucciones por ciclo. Desde el inicio de la pasada década este proceso se ha visto condicionado por la emergencia de problemas tecnológicos nuevos que han motivado una reorientación de los nuevos diseños. La conferencia tratará de revisar estos problemas y analizar las soluciones microarquitectónicas adoptadas motivando la introducción de los diseños multicore y multithreading actuales e introduciendo los aspectos arquitectónico y de programación de estas nuevas arquitecturas.



El Prof. Francisco Tirado es actualmente catedrático de Arquitectura y Tecnología de

Computadores en la Universidad Complutense de Madrid, cargo que ha ocupado desde 1986. Ha trabajado en diferentes áreas de investigación dentro de la arquitectura de sistemas, computación de altas prestaciones, computación en GRID, diseño automático de CI y arquitectura del procesador. Dentro de la UCM, ha impulsado la creación de un grupo de investigación (Artecs) que coordina en la actualidad, con más de 30 investigadores que tienen un reconocido prestigio en los campos de la Computación de Altas Prestaciones y la Síntesis de Sistemas Digitales.

Hasta la actualidad es coautor de unas 280 publicaciones en revistas y conferencias internacionales.

2. Multiphase reacting flows : from modelling issues to high performance computing ***Marc Massot***

AULA MAGNA FP-UNA, VIERNES 22 DE OCTUBRE, 18:30 HS.

The modeling and simulation of multiphase reacting flows covers a large spectrum of applications ranging from combustion in automobile and aeronautical engines to atmospheric pollution, non-linear chemical dynamics as well as biomedical engineering. In the framework of this talk, we will mainly focus on a disperse liquid phase carried by a gaseous flow field which can be either laminar or turbulent; however, this spray can be polydisperse, that is constituted of droplets with a large size spectrum. The related gaseous phase chemistry can potentially involve complex mechanism involving many species and the reaction fronts usually present steep gradients and very small localized spatial scales. Thus, such flows involve a large range of temporal and spatial scales which have to be resolved in order to capture the dynamics of the phenomena and provide reliable and eventually predictive simulation tools. Even if the power of the computer resources regularly increases, such very stiff problems can lead to serious numerical difficulties and prevent efficient multi-dimensional simulations. The purpose of this talk is to show that all the necessary steps in order to develop a new generation of computational code have to be designed at the same time with a high level of coherence: mathematical modeling, development of new dedicated numerical methods, implementation of optimized algorithms as well as validations of both model and methods using either experimental measurements or other codes. By the way, we will introduce both a new class of models for the direct numerical simulation of spray dynamics, as well as a set of dedicated numerical methods and prove that such an approach has the ability, once validated, to lead to high performance computing on parallel architectures. Recent extension of the original method will be provided and this will also be the occasion to present very recent results obtained during the Summer Program 2010 of the Center for Turbulence Research at Stanford University.

El prof. Massot trabaja actualmente en la *École Centrale Paris* de Francia en el campo de Mecánica de Fluidos y Matemática Aplicada y es investigador en el laboratorio EM2C del CNRS (*Centre National de la Recherche*

Scientifique) de Francia. Realizó sus estudios de doctorado en la Ecole Polytechnique en Matemáticas Aplicadas; fue autor y colaboró en varias publicaciones a nivel internacional. Sus campos de investigación actuales incluyen el modelado matemático, el análisis matemático, el análisis numérico y la computación científica y de alto desempeño.

3. Investigación Operativa y la Toma de Decisiones – Una experiencia

Nelson Maculan

AUDITORIO CENTRAL DE LA UNA, MARTES 19 DE OCTUBRE, 18:30-19:30 HS.



Actualmente enseña en la Universidad Federal de Río de Janeiro, el el área de Ingeniería de Sistemas y Ciencias de la Computación. Sus áreas de investigación incluyen programación matemática, investigación de operaciones, ciencias de la computación, análisis numérico, matemáticas aplicadas y bioinformática. Ha escrito y colaborado en diversos libros y papers sobre estos temas. Así también, ha dado conferencias en países a lo largo de todo el mundo.

4. Scalable Distributed Web Search Engines

Mauricio Marin

AUDITORIO CENTRAL DE LA UNA, MARTES 19 DE OCTUBRE, 19:30-20:30 HS.

Dealing efficiently with multiple user queries, each at a different stage of execution at any given instant of time, is a key issue in large-scale Web search engines. Here the use of suitable algorithms and heuristics devised to achieve a high query throughput upon the least possible amount of hardware, and yet remain stable under sudden peaks in query traffic, is critical to efficient data center operation. Achieving this goal is quite beyond the possibilities of a single indexing data structure and respective query processing algorithm. Indeed, current practice in Web search engines clearly indicates that such a goal is only feasible through a combination of indexing, caching and parallelism, all devised to work together so that they -- as a whole -- lead to efficient and scalable performance upon well-dimensioned hardware. This talk describes combinations of such strategies and shows that they lead to Web search engines devised to be efficient in terms of query throughput, individual query latency, and power consumption.



Dr. Mauricio Marin is an associate researcher at Yahoo! Research Santiago and a full-professor at the Department of Informatics Engineering of the University of Santiago of Chile. He holds a PhD in Computer Science from University of Oxford, UK, and a MSc from

University of Chile. He is a former full professor at the University of Magallanes, Chile. His research work is on parallel computing and distributed systems with applications in Web search engines and crawling. He has obtained research grants on parallel query processing upon distributed metric-space databases and scalable parallel algorithms and data structures for text search and indexing.

5. Liquid Journals: how social computing and liquid knowledge will shape and navigate information waters

Fabio Casati

AULA MAGNA FP-UNA, JUEVES 21 DE OCTUBRE, 18:30 HS.

The way scientific knowledge is created, disseminated, searched, consumed, and evaluated has been essentially the same for centuries, largely oblivious to the Web revolution. Despite technology, today it is very hard to find knowledge of interest - probably hardest than it is to find interesting web pages - despite a huge effort done by the community in filtering it (e.g., via peer review) and in evaluating it (via metrics of dubious effectiveness, such as citation count). The growing community of researchers and exploding number of publication venues makes it even more difficult to navigate the sea of information.

Liquid journals are a way to find, consume, create, and share interesting and relevant scientific knowledge.

They are based on a few key intuitions. The first is that scientific knowledge is not communicated (only) via isolated scientific papers, linked to each other via citations. Rather, it exists in different kinds (data, experiments, ideas,...), it is communicated in different forms (papers, talks, blogs), it evolves almost continuously over time, and it is connected in a knowledge network (papers describe experiments that are built over datasets, all based on ideas inspired from other talks or blogs). The second intuition is that the network which would be so useful to navigate in the sea of scientific knowledge is not objective, but is rather subjective. Whether a paper is inspired from another, or whether a person contributed to a paper may be proven facts or may be opinions, which may also take different forms. The third intuition is that we can use the power of the community as editors to help us select knowledge among a sea of information, rather than leaving this role to a selected few. The fourth intuition is that editors and the community of readers can create knowledge. This is a huge potential that is currently untapped for various social and technical reasons, but that can be used to service the scientific community.

This talk will present the main ideas behind liquid journals and then focus on the services and social computing side, describing the IT infrastructure that enables each of us to provide, create, consume, and share scientific knowledge. We will see that the "social" aspect lies both in creating IT services that facilitate the social content creation and editing effort as well as in having the community actually contribute to building the very same IT services that support this.



Fabio Casati is Professor of Computer Science at the University of Trento. He got his PhD from the Politecnico di Milano and then worked for over 7 years in Hewlett-Packard USA, where he was technical lead for the research program on business process intelligence. Fabio has also contributed (as software and data architect) to the development of several HP commercial products and solutions in the area of web services and business process management. In Trento, he is leading or participating to several FP7 projects, is active in many industry-funded projects, both local and international, and has over 20 patents. His passions are now in social informatics, or, informatics at the service of the community. His latest efforts are on prevention of non-communicable diseases, on remote and real-time healthcare, on collaborative programming, and on models for scientific disseminations that can help scientists work in a more efficient way.

6. SOFTWARE LIBRE AND THE FUTURE OF LEARNING

Walter Bender

AUDITORIO CENTRAL DE LA UNA, LUNES 18 DE OCTUBRE, 18:30 HS.



Walter Bender is a senior research scientist and director of the Electronic Publishing group; he also directed the Gray Matters special interest group, which focused on technology's impact on the aging population. In 1992, Bender founded the News in the Future consortium and has been a member of the Lab's SIMPLICITY, Things That Think, and Digital Life consortia. Bender joined the Architecture Machine Group at MIT in 1978, after receiving his BA from Harvard University in 1977. He received his MS at MIT in 1980. A founding member of the Media Laboratory, throughout his career Bender has engaged in the study of new information technologies, particularly those that affect people directly. Much of the research addresses the idea of building upon the interactive styles associated with existing media and extending them into domains where a computer is incorporated into the interaction. He has participated in much of the pioneering research in the field of electronic publishing and personalized interactive multimedia. Before taking his leave of absence from MIT, Bender was executive director of the MIT Media Laboratory, and holder of the Alexander W. Dreyfoos Chair. While on leave, Bender served as president for software and content development of One Laptop per Child, the not-for-profit association that is developing and deploying technologies that will revolutionize how the world's children engage in learning. He currently heads Sugar Labs, and continues

work on the award-winning Sugar Learning Platform. They recently announced "Sugar on a Stick," which gives children access to Sugar on any computer by using a USB key.