

WEDNESDAY 24 AUGUST

Living in a Heterogenous World: How scientific workflows help automate science and what we can do better?

Ewa Deelman, Research Professor and Research Director, USC Information Sciences Institute, USA



Ewa Deelman
Research Professor and
Research Director
USC Information Sciences
Institute, USA

Scientific workflows are now a common tool used by domain scientists in a number of disciplines. They are appealing because they enable users to think at high level of abstraction, composing complex applications from individual application components. Workflow management systems (WMSs), such as Pegasus (<http://pegasus.isi.edu>) automate the process of executing these workflows on modern cyberinfrastructure. They take these high-level, resource-independent descriptions and map them onto the available heterogeneous resources: campus clusters, high-performance computing resources, high-throughput resources, clouds, and the edge.

WMSs can select the appropriate resources based on their architecture, availability of key software, performance, reliability, availability of cycles, storage space, among others. With the help of compiler-inspired algorithms, they can determine what data to save during execution, and which are no longer needed. Similarly to compiler solutions, they can generate an executable workflow that is tailored to the target execution environment, taking into account reliability, scalability, and performance. WMS use workflow execution engines to run the executable workflows on the target resources providing scalability and reliability.

This talk will describe the key concepts used in the Pegasus WMS to help automate the execution of workflows in distributed and heterogeneous environments. It will explore potential use of artificial intelligence and machine learning approaches to enhance automation. The talk will also help identify challenges that exist in adopting novel approaches for science at the technological and social levels

THURSDAY 25 AUGUST

Effective Congestion Management for Large-Scale Datacenters

José Duato, Full Professor, Polytechnic University of Valencia, Spain



Jose Duato
Full Professor
Polytechnic University of
Valencia, Spain

Datacenters are essential for providing Internet services. As the number of client requests per time unit and their complexity keep increasing, datacenters are adopting computing solutions to scale with the demand, and provide appropriate support for interactive services. In particular, computing accelerators (mostly GPUs, but also TPUs, FPGAs, etc) have become very popular, and some recent designs even incorporate network ports in those devices to directly attach them to the interconnection network. As system size increases, the cost of the interconnection network grows faster than system size, thus becoming increasingly important to carefully design it to prevent over-provisioning. However, by doing so, the network operation point moves closer to saturation and sudden traffic bursts may lead to congestion. This situation is aggravated by the recent introduction of flow control in datacenter networks to cope with RDMA requirements, and network power management. The result is massive performance degradation whenever some network region becomes congested. Moreover, performance degradation may remain for long even after the traffic bursts that congested the network have already been transmitted.

This keynote will show why congestion appears in an interconnection network, how it propagates, and why performance may degrade so dramatically. Different kinds of congestion will be identified. Also, a global solution to effectively address the congestion problem will be proposed. It consists of several complementary mechanisms that accurately identify the congestion sources and cooperate to address all kinds of congestion, operating at different time scales. Some of these mechanisms have been recently incorporated into commercial products and are being standardized.

FRIDAY 26 WEDNESDAY

Programming Big Data Analysis: Towards Data-Centric Exascale Computing

Domenico Talia, Professor of Computer Engineering, DIMES of University of Calabria, Italy



Domenico Talia
Professor of Computer
Engineering
DIMES of University of
Calabria, Italy

Software applications today are strongly data driven. For this reason programming models and tools and novel architectures have been recently studied and developed to extract valuable information from Big Data, addressing data complexity, scalability, and/or high velocity. Analytics and machine learning on Big Data sources are not feasible through sequential algorithms to obtain in a reasonable time models and patterns from huge volumes of data. For this reason, parallel computers, such as many and multi-core systems, Clouds, and multi-clusters, along with parallel and decentralized algorithms and systems are required to analyze Big Data sources and repositories. In this direction Exascale computing systems represent the next step. Exascale systems refer to high performance computing systems capable of at least one exaFLOPS, so their implementation is representing a very significant research and technology move. In fact, cluster computers and Cloud platforms used today can store very large amounts of data, however they do not provide the high performance expected from massively parallel Exascale systems. This is the main motivation for developing Exascale platforms that will represent the most advanced model of supercomputers.

Data analysis solutions advance by exploiting the power of data mining and machine learning techniques and are changing several scientific and industrial areas. Therefore, it is vital to design scalable solutions for processing and analysis such massive datasets. Scalability and performance requirements are challenging conventional data storages, file systems and database management systems. Architectures of such systems have reached limits in handling very large processing tasks involving petabytes of data because they have not been built for scaling after a given threshold. This condition claims for new hardware architectures and data analysis software solutions that must process Big Data for extracting complex predictive and descriptive models. To reach Exascale size, it is in fact required to define new programming models and languages that combine abstraction with both scalability and performance Hybrid models (shared/distributed memory) and communication mechanisms based on locality and grouping are currently designed as promising approaches. Parallel applications running on Exascale systems require to control millions of threads running on a very large set of cores. Such applications need to avoid or limit synchronization, use less communication and remote memory, and handle with software and hardware faults that could occur.



if you are joining us in @UofGlasgow for @euro_par this August please check out our website where you can find information on Travel and Activities
<https://t.co/ccQ7a5KmGx>

22.06.2022 - 13:09

The Famous Edinburgh Festival @edfest will be happening between the 5th and 28th of August. If you are joining as for @euro_par in Glasgow please check it out. Edinburgh is a quick train ride away!

14.06.2022 - 14:33

@euro_par announced the PhD symposium for 2022 please find more information here: <https://t.co/z4g8OaIZGO> * Key Dates Submissions: 20 May 2022 Notifications: 27 June 2022

27.04.2022 - 16:37

Euro-Par Achievement Award and Keynote Speakers for Euro-Par'22: Euro-Par Achievement Award: Ewa Deelman <https://t.co/BUrn02fZ>; Keynote Speakers: José Duato <https://t.co/TTQXKpzy>; Domenico Talia <https://t.co/BvimOh02Lv>; Ewa Deelman <https://t.co/BUrn02fZe>

27.04.2022 -



GOLD SPONSORS



SILVER SPONSORS



This site uses third-party website tracking technologies to provide and continually improve our services, and to display advertisements according to users' interests. I agree and may revoke or change my consent at any time with effect for the future.

DENY

ACCEPT

MORE

This site uses third-party website tracking technologies to provide and continually improve our services, and to display advertisements according to users' interests. I agree and may revoke or change my consent at any time with effect for the future.

DENY

ACCEPT

MORE

Powered by  usercentrics &

