



Integrating the access to persistent data with parallel programming



Toni Cortes,
Associate Professor
Universitat Politècnica de Catalunya &
Barcelona Supercomputing Center

Περίληψη – Abstract

Since the beginning, persistent data and non persistent data have been treated as two separate abstractions. A clear example is that the model used to store data into volatile memory (mainly objects and their relations) is completely different from the model used to store the same data into a persistent storage (mainly tables or files). This differentiation between data has many negative side effects because persistent data cannot be integrated into the programming model. This lack of integration causes, among others, the following problems i) moving computation to the data becomes a complex task (deployment can become an arduous task) , ii) the extraction of potential data parallelism by the programming model is very difficult (the programming model is unaware of where the data really is), and iii) offering a mechanisms to really share data without taking the control from the data owner becomes nearly impossible (we will show that today data is not really shared).

In this talk, we will present dataClay, a new-generation object storage and its integration with the COMPSs programming model. This new way to handle data (and code), and its perfect fit with a parallel programming model will eliminate all the afore-mentioned problems easing the task of implementing data-centric programs while full advantage of the available parallelism.

[Toni Cortes](#) is the manager of the storage-system group at the BSC (since 2006) and is also an associate professor at Universitat Politècnica de Catalunya (since 1998). He received his M.S. in computer science in 1992 and his Ph.D. also in computer science in 1997 (both at Universitat Politècnica de Catalunya).

Since 1992, Toni as been teaching operating system and computer architecture courses at the Barcelona school of informatics (UPC) and from 2000 to 2004 he also served as vicedean for international affair at the same school.

His research concentrates in storage systems, programming models for scalable distributed systems and operating systems. He has published 103 technical papers (25 journal papers and 78 international conferences and workshops), 2 book chapters, and has co-edited one book on mass storage systems. In addition, he has also advised 10 PhD thesis.

Dr. Cortes has been involved in several EU projects (Paros, Nanos, POP, XtremOS, Scalus, IOlanes, PRACE, MontBlanc, EUDAT, Big Storage, IOstack, NextGenIO, Rethinkbig, and Severo Ochoa) and has also participated in cooperation with IBM (TJW research lab) on scalability issues both for MPI and UPC. He is also editor of the Cluster Computing Journal and the coordinator of the SSI task in the IEEE TCSS. He has served in many international conference program committees and/or organizing committees and was general chair for the Cluster 2006 conference, LaSCO 2008, XtremOS summit 2009, and SNAPI 2010. He is also served as the chair of the steering committee for the Cluster conference series (2011-2014). His involvement in IEEE CS has been awarded by the "Certificate of appreciation" in 2007.

Τρίτη 24/05/2016 – 11:00

Αίθουσα Σεμιναρίων,

**Κτίριο Μηχανικών Η/Υ & Πληροφορικής
Πανεπιστήμιο Ιωαννίνων**