



Andrei Tchernykh is a researcher in the Computer Science Department, CICESE Research Center, Ensenada, Baja California, Mexico. From 1975 to 1990 he was with the Institute of Precision Mechanics and Computer Engineering of the Russian Academy of Sciences (IPMCE, Moscow, Russia). He received the Ph.D. in Computer Science from IPMCE in 1986. In CICESE, he is a coordinator of the Parallel Computing Laboratory. He is member of the National System of Researchers of Mexico (SNI), Level II. He leads number of national and international research projects. He served as a program committee member and organizer of professional international conferences. His main interests include scheduling, load balancing, adaptive resource allocation, scalable energy-aware algorithms, multi-objective optimization, heuristics, meta-heuristics, and incomplete information processing (<http://usuario.cicese.mx/~chernykh/>).

Understanding uncertainty in Cloud Computing

Andrei Tchernykh

Computer Science Department, CICESE, Ensenada, BC, México. chernykh@cicese.mx, +521(646)1608350

Clouds differ from previous computing environments in the way that they introduce a continuous uncertainty into the computational process. The uncertainty becomes the main hassle of cloud computing bringing additional challenges to both end-users and resource providers. Both require waiving habitual computing paradigms, adapting current computing models to this evolution, and designing novel resource management strategies to handle uncertainty in an effective way.

In this talk, we discuss the role of uncertainty in the resource/service provisioning, investment, operational cost, programming models, etc. that have not yet been adequately addressed in the scientific literature. We discuss several major sources of uncertainty: dynamic elasticity, dynamic performance changing, virtualization, loosely coupling application to the infrastructure, among many others. A workload in such an environment is not predictable and can be changed dramatically. It is impossible to get exact knowledge about the system. Parameters such as an effective processor speed, number of available processors, and actual bandwidth are changing over the time. Elastic escalation process has a higher repercussion on the QoS, but adds another factor of uncertainty.

In most existing solutions, it is assumed that VMs are predictable and stable in their performance. In actual cloud infrastructures these assumptions do not hold. Providers might not know the quantity of data and computation required by users. For example, every time when a user requires a status of his e-mail or bank account, it could generate different amount of data and take different time for delivering. A pool of virtualized, dynamically scalable computing resources, storages, software, and services add a new dimension to the problem. The manner in which the service provisioning can be done depends not only on the service property and needed resources, but also users that share resources at the same time, in contrast to dedicated resources governed by a queuing system.