

BSC releases COMPSs version 2.2 at SC

The Workflows and Distributed Computing team at the Barcelona Supercomputing Center is proud to announce a new release, version 2.2 (codename Camellia), of the programming environment COMPSs.

- The Barcelona Supercomputing Center offers to the HPC community a set of tools that helps developers to program and execute their applications efficiently on distributed computational infrastructures
- This release includes new features that enable the existence of computation and dataflow tasks in a single application through the integration with the Decaf system; extensions to the runtime to support elasticity in Slurm-managed clusters and Singularity containers; and an enhanced support for C/C++ applications with a new persistent worker implementation and a worker software-cache.
- Attend a PyCOMPSs tutorial at the Intel HPC Developer Conference, paper presentations at PyHPC and WACCPD workshops and see life demos at BSC booth #1975.

This version of COMPSs, available from today, updates the result of the team's work in the last years on the provision of a set of tools that helps developers to program and execute their applications efficiently on distributed computational infrastructures such as clusters, grids and clouds. COMPSs is a task based programming model known for notably improving the performance of large scale applications by automatically parallelizing their execution.

COMPSs has been available for the last years to the MareNostrum supercomputer users and to the Spanish Supercomputing Network and has been adopted in several research projects such as OPTIMIS, VENUS-C, EUBrazil OpenBio, EUBrazil CloudConnect, transplant, EGI and ASCETIC. In these projects COMPSs has been applied to implement use cases provided by different communities across diverse disciplines as biomedicine, engineering, biodiversity, chemistry, astrophysics and earth sciences. Currently it is also under extension and adoption in applications in the projects EUBRA BIGSEA, NEXTGenIO, MUG, mf2C, TANGO, the CoE BioExcel, the EXPERTISE ETN, and the Human Brain Project flagship.

The new release comes with support to dataflow systems, through the integration with Decaf, a dataflow system for the coupling of tasks in a workflow. The dataflow can perform arbitrary data transformations ranging from simply forwarding data to a complex parallel program executing the required transformation. The integration performed in PyCOMPSs, the Python binding of COMPSs, offers to the user the possibility of developing PyCOMPSs workflows with tasks that are Decaf pipelines. Users of Decaf and PyCOMPSs can benefit from this integration in the following way. On one hand, PyCOMPSs users can easily integrate Decaf pipelines in their applications for efficient in situ computations. On the other hand, Decaf users can easily extend their applications to distributed locations when data are located in different sites or when a single location is insufficient for performing the required computation.

COMPSs runtime supports elasticity but since now this feature has been constrained to cloud environments. However, with the current support of Slurm to elasticity in clusters, COMPSs has been extended to interact with Slurm in such a way that COMPSs applications can benefit from this feature when running in clusters managed by Slurm. This feature has also been extended to clusters managed with Singularity containers.

Recent releases of COMPSs included a persistent worker for Java and Python, but this feature was missing for C/C++. COMPSs release 2.2 not only includes this feature, which improves the performance of the applications due to a smaller overhead on task execution, but also comes with a worker software-cache which maintains in the worker memory the tasks' arguments in order to exploit its potential locality.

COMPSs 2.2 comes with other minor new features, extensions and bug fixes

At SC17, PyCOMPSs/COMPSs will have a large presence, with tutorials, presentations and demos. A tutorial presentation with hands-on at the Intel HPC Developer Conference will be delivered on Saturday 11th Nov. A presentation of PyCOMPSs and its application to linear algebra kernels will be delivered on PyHPC workshop, and a presentation of the GPU support with COMPSs-mobile will be delivered in the WACCPD workshop. Furthermore, live demos on the use of PyCOMPSs with the Jupyter notebook and on real-use cases that integrate MPI simulations, new storage technologies and parallelized queries with PyCOMPSs.

COMPSs has had around 1000 downloads last year and is used by around 20 groups in real applications. COMPSs has recently attracted interest from areas such as image recognition, genomics and biodiversity, where specific courses and dissemination actions have been performed.

The packages and the complete list of features are available in the Downloads page. A virtual appliance is also available to test the functionalities of COMPSs through a step-by-step tutorial that guides the user to develop and execute a set of example applications.

Additionally, a user guide and papers published in relevant conferences and journals are available.

For more information on COMPSs please visit our [webpage](#).

About Barcelona Supercomputing Center

Barcelona Supercomputing Center (BSC) is the national supercomputing centre in Spain. BSC specialises in High Performance Computing (HPC) and its mission is two-fold: to provide infrastructure and supercomputing services to European scientists, and to generate knowledge and technology to transfer to business and society.

BSC is a Severo Ochoa Center of Excellence and a first level hosting member of the European research infrastructure PRACE (Partnership for Advanced Computing in Europe). BSC also manages the Spanish Supercomputing Network (RES).

The Workflow and Distributed Computing team at the Barcelona Supercomputing Center aims to offer tools and mechanisms that enable the sharing, selection, and aggregation of a wide variety of geographically distributed computational resources in a transparent way. The research done in this team is based in the former expertise of the group, and extending it towards the aspects of distributed computing that can benefit from this expertise. The team at BSC has a strong focus on programming models and resource management and scheduling in distributed computing environments.

For more info, please contact the communication [at] bsc [dot] es (communication department) at BSC.

Tel: +34 620 429 956 (Gemma Ribas)

Barcelona Supercomputing Center - Centro Nacional de Supercomputación

Source URL (retrieved on 14 ago 2024 - 05:53): <https://www.bsc.es/ca/news/bsc-news/bsc-releases-compss-version-22-sc>