6th OpenFOAM Workshop
PennState University, USA
13-16 June, 2011

Constellation Platform - a distributed super-computer enabling parallel applications on volunteer PCs for aerospace science

Dipl.-Ing. (FH) Andreas HORNIG[*‡1] and Dr. Jaspal SUBHLOK[2]

[1] Constellation Platform, Stuttgart, Germany
[2] Volpex Group, University of Houston, Texas, USA

February 15, 2011

Distributed Computing, Parallel Computing, On Volunteer PCs, Via Internet Connected Nodes

CONSTELLATION:

Constellation is a platform for different aerospace related projects that need intensive computational power. The platform supports the efforts of participating projects by providing Distributed Computation capability using BOINC[1]. Constellation will send work-units of attached projects to volunteering, idle PCs where the units are progressed. The combined power of all volunteering users will help to solve important scientific tasks in fields from astronomy to aerospace-engineering beginning from student up to university projects.

In particular in aerospace sciences simulation and optimization are major tasks to minimize structure weight, to maximize thrust or to determine system reliability. These numerical tasks need a high level of resources and hardware support and Constellation platform offers distributed computing power to projects from professionals to students to solve their problems in an adequate time without the need to maintain their own super-computer or cluster system. In this way it is possible to get access to the needed computing power by a wider range of researchers who wouldn't have been able to create those resources because of financial, administration and operation or because of bureaucracy reasons.

Currently the BOINC infrastructure Constellation uses supports independent workunits. The above applications respect this limit by using splitted tasks as workunits that doesn't rely on other workunits' result. But there are aerospace tasks, like in computational fluid dynamics (CFD), that are only feasible when the complete tasks is solved in parallel to finish in a decent time. Therefore Constellation is working to extend the system to combine the advantages of distributed computing system with that of parallel clusters to create a virtual cluster where nodes are connected via the internet. We work together with Volpex[5] Group at University of Houston to bring parallel execution to Constellation and to other BOINC-projects, and we examine MAGE[2] of University of Marburg as an additional candidate.

‡ Corresponding Author: Andreas Hornig, Constellation Platform (constellation@aerospaceresearch.net)

We want to use parallel execution for CFD analysis. We want to use OpenFOAM[3] because it's used in the scientific and academic community, it offers a wide range of solvers and the main reason is that its GNU General Public License allows the use in a distributed environment. In comparison Ansys' powerful CFX and FLUENT[4] programmes use a proprietary software license that makes it difficult to use outside the given license agreement. For Constellation a CFD application is intended to be used in aerodynamics and stability simulation, in engine combustion analysis for micro- and jet engines, in ramjet, scramjet and pulse detonation engine and rocket engines and motors and many more. We want to open up this important field of parallel execution for distributed computing.

VOLPEX:

Volpex[5-7] has the goal of enabling robust execution of communicating parallel applications on volunteer PCs. The project developed the fundamental concept of autonomous redundant processes: a logical process may have multiple distributed instances that are unaware of each other, process instances can be replicated, checkpointed or recreated independently, and individual process instances can fail without application failure. The overall application progress is dictated by the state of the fastest active instance of each process. A single replication framework can potentially provide protection from process/node and network failures, varying processing speeds, and Byzantine failures caused by software and hardware errors and malicious hosts. Two prototype programming frameworks were successfully developed and validated in this project:

- Volpex Dataspace API: An abstract ``dataspace'' with asynchronous anonymous Put/Get communication operations from autonomous redundant processes.
- Volpex MPI: A highly failure resistant MPI implementation based on sender based logging.

A Volpex prototype has been validated and evaluated with NAS and other benchmarks, a Map-Reduce framework, and a Replica Exchange Molecular Dynamics code, and interfaces with BOINC[1].

REFERENCES

[1] BOINC – Berkeley Open Interface for Network Computing - http://boinc.berkeley.edu

[2] MAGE – The Marburg Ad-hoc Grid Environment - http://mage.uni-marburg.de

[3] OpenFOAM – Open Field Operation and Manipulation - http://www.openfoam.com

[4] Ansys – CFX, Fluent - http://ansys.com

[5] VOLPEX - Parallel Execution in Volunteer environment - http://volpex.cs.uh.edu/VCP

[6] 'A Robust Communication Framework for Parallel Execution on Volunteer PC Grids', By Eshwar Rohit, Hien Nguyen, Nagarajan Kanna, Jaspal Subhlok, Edgar Gabriel, Qian Wang, Margaret S. Cheung, and David Anderson, The 11th IEEE/ACM International Symposium on Cluster, Cloud, and Grid Computing , Newport Beach, CA, May 2011.

[7] VolpexMPI: an MPI Library for Execution of Parallel Applications on Volatile Nodes, By T. LeBlanc, R. Anand, E. Gabriel, and J. Subhlok. In Proc. The 16th EuroPVM/MPI 2009 Conference, Espoo, Finland, 2009. Published as Springer-Verlag LNCS,5759.