

Adaptive Execution of Parallel Programs on Grids and Clouds

Jaroslav Slawinski¹ and Vaidy Sunderam^{1*}

¹*Department of Math and Computer Science
Emory University, Atlanta, GA 30322, USA
vss@emory.edu*

Abstract. Streamlined switching between computational resources in order to select the most suitable computational environment for parallel application execution is a crucial component of utility-like computing. However, machine heterogeneity obstructs multi-target deployment for complex and multi-dependency scientific parallel codes and makes this aim intractable. We describe a proposal for a metadeployment toolkit, called ADAPT, based on reusable recipes addressing appropriate match-up between an application and an execution platform. Our research aims at exploring challenges posed by transparent application deployment with all its prerequisites on heterogeneous resources. As some IaaS clouds and grids accept customized OS images, we explore application-oriented image assembly to further improve deployment for these specific targets. We explain how our approach increases “usability” of various resources for parallel applications and simplifies arcane build processes.

Keywords: Cloud Computing, Computational Grids, High Performance Computing, Software Deployment, Build

1. Introduction

We examine High Performance Computing applications manifested as parallel programs with the following attributes: source code availability, nontrivial software dependencies, and target optimization requirements. These applications typically execute on large parallel machines. On the other hand, the recent surge in on-demand computational offerings encourages users to experiment with parallel applications on a broader assortment of resources. Despite unavoidable performance degradation, new platforms attract for various reasons: instant resource availability (no job queues; any number of hosts), cost (no upfront expenses), or greater control (root privileges). However, to facilitate the execution on other resources, an improved, automatic approach to parallel application deployment is required as end-users cannot rely on user support offered at HPC centers.

This paper delivers a proposal and a use case of a system that provides automatic, adaptive, and transparent multi-targeted deployment solution. This design employs reusable deployment recipes that encapsulate and store expert knowledge related to software conditioning. Proper chaining of these recipes formulate automatic deployment scripts that probe and soft-condition the target environment until the considered application is successfully installed. As a result, users may focus on executing their parallel applications on wider range of resources without obstacles related to deployment phases. Our proposal contributes to usability enhancement of computational offerings and has the potential to

*Corresponding Author: Vaidy Sunderam, Math & CS, Emory University, Atlanta, GA 30322. Email: vss@emory.edu
Research supported in part by Air Force Office of Scientific Research grant FA-9550-15-1-0950