

Abstract

Performance Analysis of Virtual vs. Physical Machines for Automated Judging in Programming Contests

John Clevenger and Cody Jackson
California State University, Sacramento

Although many programming contests use human judges to evaluate team submissions, more and more contests are for a variety of reasons relying on automated “computer judging”. This is particularly true in light of the fact that most of the widely-available Contest Control Systems provide direct support for automated judging. In most such systems this automated judging is handled by allocating a separate set of machines called something like “auto-judges” (AJ’s) to perform the judging operations.

In order to insure a high-quality experience for teams, each team is normally given their own physical machine. The need for additional machines for other functions, such as judging, sometimes makes the allocation of computing resources a problem. This is for example an ongoing problem at the ICPC World Finals (WF), where a substantial number of physical AJ machines are allocated for judging purposes.

One potential solution to the resource allocation problem is to assign the AJ function to virtual machines. This however poses a new dilemma: how to insure that the execution time of a team’s submission does not vary too widely between the team’s physical machine and the (virtualized) AJ machine. This presentation describes ongoing work designed to characterize and measure the extent of this problem as it relates to programming contests, with specific attention to the ICPC World Finals environment.

We present a newly-developed framework, the *ICPC Performance Framework (IPF)*, which has been constructed to examine the differences between the execution of team submissions in physical and virtualized AJ environments. We have collected together tens of thousands of actual team submissions from numerous past ICPC World Finals, and run each of these submissions through our IPF thousands of times using a variety of both physical and virtual auto-judges. Analyzing the runtime data from this large set of actual contest submissions has allowed us to develop some preliminary conclusions regarding the potential ramifications of switching a contest such as the ICPC World Finals from a physical to a virtualized judging environment.

In this presentation we describe the organization of the IPF implementation, examine the preliminary results of the analysis of the data generated using ICPC World Finals contest submissions, and present some conclusions and recommendations regarding issues which need to be considered when considering switching from physical to virtual machines for automated judging.