Evaluating the Advantage of Reactive MPI-aware Power Control Policies

Daniele Cesarini[1], Carlo Cavazzoni[2], Andrea Bartolini[1]
[1]University of Bologna, viale del Risorgimento, 2, 40136 Bologna, Italy
[2]Cineca, Via Magnanelli, 6/3, 40033 Casalecchio di Reno, Italy
{daniele.cesarini,a.bartolini}@unibo.it
c.cavazzoni@cineca.it

Power consumed is an essential factor that worsens the performance and costs of today and future supercomputer installations. In state of the art (SoA), some approaches have been proposed to reduce the energy consumed by applications running in this machine by reducing the operating frequency of the computational engine during communication regions in MPI parallel applications. Algorithms in the SoA rely on the capability of predicting at run-time the duration of these communication regions before their execution. The COUNTDOWN approach tries to do the same by mean of a purely reactive timer based policy. In this paper, we compare the COUNTDOWN algorithm with SoA predictive-based algorithm, showing that timer based policies are more effective in extract power saving opportunities, saving power inducing a lower overhead. When running in a Tier1 system COUNTDOWN achieves 5% more energy saving with lower overhead than state-of-the-art proactive policy. This suggests that reactive policies are more suited then proactive approaches for communication aware power management algorithms.

Keywords: HPC, MPI, power management, reactive policy, DVFS, NPB, energy efficiency, parallel programming.