Energy-Efficient Job-Assignment Policy with Asymptotically Guaranteed Performance Deviation

Jing Fu[1], Bill Moran[2]
1 School of Mathematics and Statistics, the University of Melbourne, Australia
2 Department of Electrical and Electronic Engineering, the University of Melbourne, Australia
Email: jing.fu@unimelb.edu.au; wmoran@unimelb.edu.au

We study a job-assignment problem in a large-scale server farm system with geographically deployed servers as abstracted computer components (e.g., storage, network links, and processors) that are potentially diverse. We aim to maximize the energy efficiency of the entire system by effectively controlling carried load on networked servers. A scalable, near-optimal job-assignment policy is proposed. The optimality is gauged as, roughly speaking, energy cost per job. Our key result is an upper bound on the deviation between the proposed policy and the asymptotically optimal energy efficiency, when job sizes are exponentially distributed and blocking probabilities are positive. Relying on Whittle relaxation and the asymptotic optimality theorem of Weber and Weiss, this bound is shown to decrease exponentially as the number of servers and the arrival rates of jobs increase arbitrarily and in proportion. In consequence, the proposed policy is asymptotically optimal and, more importantly, approaches asymptotic optimality quickly (exponentially). This suggests that the proposed policy is close to optimal even for relatively small systems (and indeed any larger systems), and this is consistent with the results of our simulations. Simulations indicate that the policy is effective, and robust to variations in job-size distributions.

Server farm vendors deploy a variety of computer components, such as CPUs and disks, to meet various types of inquiries from Internet users, and different generations of these components are present simultaneously because of partial upgrading of old components and purchasing new ones over time. A diversity of physical computing/storage components are available for use in cloud computing platforms, and are abstracted (virtualized) as resources with varying attributes. All of these have resulted in heterogeneity as an important feature in attempting to undertake research on server farms.

On the other hand, large modern server farms with hundreds of thousands of computer components (abstracted servers) require all scheduling policies to be scalable. Existing job-assignment policies for network resource allocation problems applicable for

practical scenarios with heterogeneous servers and jobs, were studied as static optimizations. Profits to be gained through dynamic release and reuse of resources were ignored. Here we use methods of stochastic optimization that capture dynamic properties of a system.

To maximize the energy efficiency, defined as the ratio of the long-run average throughput to the long-run average power consumption, in a stochastic system with heterogeneous servers, an asymptotically optimal job-assignment policy can be proposed based on *restless multi-armed bandit* techniques. This optimal policy approaches the optimal solution as the numbers of servers in different server groups tend to infinity proportionately. Nonetheless, the achieved asymptotic optimality in the past work was restricted in two aspects: a) although modern server farms are normally large enough to be close to the asymptotic regime, the critical value, above which the numbers of servers are "sufficiently large", remains unclear, and b) it was assumed that any server in the server farm can serve any arriving job if it has a vacant slot in its buffer. This constraint is not appropriate for geographically separated or functionally varied computer components (abstracted servers).

We aim to maximize the energy efficiency and study the deviation between a newly proposed, scalable policy and the true optimal solution; particularly studying the relationship between this deviation and the number of servers in the system. We refer to the abstracted servers that are potentially able to serve a job as the available servers of this job.

The primary contribution of our work is a sharpening of the asymptotic optimality results in a heterogeneous server farm, discussed in the existing work in literature. Specifically, we prove that, when the job sizes are exponentially distributed and the blocking probabilities of jobs are always positive, there is a hard upper bound on the deviation between a simple, scalable policy and the optimized energy efficiency in the asymptotic regime; this upper bound diminishes exponentially as the number of servers in server groups and the arrival rates of jobs tend to infinity proportionately. In other words, the scalable policy approaches asymptotic optimality quickly (exponentially) as the size of the optimization problem increases. We refer to this upper bound as the deviation bound, and the policy as Priorities accounting for Available Servers (PAS), as it is a priority-style policy and applicable for a system with different sets of available servers.