Chapter 103
Improving Energy–Efficiency of Scientific Computing Clusters

Tapio Niemi
Helsinki Institute of Physics, Finland

Jukka Kommeri
Helsinki Institute of Physics, Finland

Ari-Pekka Hameri
University of Lausanne, Switzerland

ABSTRACT

The authors applied operations management principles on scheduling and allocation to scientific computing clusters to decrease energy consumption and to increase throughput. They challenged the traditional one job per one processor core scheduling method commonly used in scientific computing with parallel processing and bottleneck management. The authors tested the effect of increased parallelism by using different test applications related to high-energy physics computing. The test results showed that at best their methods both decreased energy consumption down to 40% and increased throughput up to 100%, compared to the standard one task per CPU core method. The trade-off is that processing times of individual tasks get longer, but in scientific computing, the overall throughput and energy-efficiency are often more important.

INTRODUCTION

Scientific computing clusters are widely used in many research fields, especially in experimental physics, astronomy, and bio-sciences. Computing intensive research easily deploys thousands, even hundreds of thousands, of CPUs to analyze various data sets and models. These clusters can be viewed as production resources processing jobs consisting of numerous tasks. The tasks can be processed by different resources, and finally the jobs are assembled together to be delivered back to the cluster customers. Operating such a cluster has its own cost structure related to capital invested, energy consumed, maintenance work, and facility related costs. In all, a computing cluster closely resembles an industrial production unit, thus the working hypothesis for this chapter is

to apply operations management principles used in manufacturing to improve computing cluster productivity and overall efficiency.

Green computing, as it is now often referred to, seeks to improve efficiency of computing centers. It is a wide topic incorporating issues like data centre locations near cheap energy sources (Brown & Reams, 2010), minimizing so called e-waste (Hanselman & Pegah, 2007), designing optimal cooling infrastructure and running the centre in an optimal way (Marwah, et al., 2009). Generally, energy and resource optimization in scientific computing has mostly focused on hardware and infrastructure issues, for example the development of more efficient hardware or optimizing cooling of computer centers, and has not focused as much on operational methods such as workload management and even less on operating systems and application software optimization for energy-efficiency.

The most well known method for comparing energy efficiency of data centers is Power Usage Effectiveness (PUE) metrics (Belacy, 2008). This is a ratio of the total facility power / IT equipment power. It indicates how much of the energy is lost in cooling, power distribution, and other infrastructures. However, it does not indicate how efficiently IT resources are operated. Limitations of PUE metrics have been recognized by the Green IT Promotion Council (2010), which proposes three additional metrics: ITEU (IT Equipment Utilization)—IT equipment usage in data center; ITEE (IT Equipment Energy Efficiency)—total rated capacity of IT equipment / total rated energy consumption of IT equipment; and GEC (Green Energy Coefficient)—Green (natural energy) energy / total energy consumption of data center. Based on these metrics it is possible to calculate Datacenter Performance Per Energy used (DPPE) metrics as follows: DPPE = (ITEU x ITEE x 1/PUE) / (1-GEC). Further, the report gives four methods for improving energy efficiency: 1) Operating the data center in an efficient way, i.e. reducing amount of hardware and increasing its utilization; 2) Installing energy efficient hardware; 3) Improving energy efficiency of non-IT infrastructure; and 4) Using renewable energy.

In this work, we focus on a typical computing problem in high-energy physics: How to process a large set of jobs efficiently. While most of the existing work on high performance computing focuses on optimizing processing time of individual computing jobs, we try to optimize energy consumption per computing job and the total processing time of the set of jobs by choosing an optimal scheduling policy.

We study this problem by using real experimental physics data, computing jobs, and dedicated computing clusters. CERN, the European Organization of Nuclear Research in Geneva, provided us with a unique possibility to experiment and test the hypothesis. The Large Hadron Collider (LHC) experiment at CERN produces about 15 petabytes of data in a year. The overall computing infrastructure comprises numerous computing clusters of alternative sizes, yet the total amount of CPUs is over 100,000 in over 140 computing centers. Efficient management of these computing resources is vital for the success of the project, which is foreseen to be active for the next 20 years.

Basically, the problem is similar to production management in any manufacturing facility. In manufacturing, this kind of optimization problem can lead to a trade-off situation: improving energy-efficiency can weaken throughput. However, within the context of the computing center our tests indicated that these two aims are not necessarily contradictory, meaning that optimizing system throughput also improves energy efficiency. Our approach is based on a conclusion that computers should run near full power or be turned off, since the fixed power consumption is around 50% of the full power of the server. This chapter is based on our earlier papers (Hameri & Niemi, 2010; Niemi, et al., 2009a, 2009b).

16 more pages are available in the full version of this document, which may be purchased using the "Add to Cart" button on the publisher's webpage: www.igi-global.com/chapter/improving-energy-efficiency-scientific-computing/69374

Related Content

New Generation of Productive Workers: How Millennials' Personal Values Impact Employee Productivity in Industry 4.0
Rok Cresnar (2021). *Research Anthology on Cross-Industry Challenges of Industry 4.0 (pp. 1632-1646).*
www.irma-international.org/chapter/new-generation-of-productive-workers/276893

Creative Regions in Europe: Exploring Creative Industry Agglomeration and the Wealth of European Regions
Blanca de-Miguel-Molina, José-Luis Hervás-Oliver, Rafael Boixand María de-Miguel-Molina (2013). *Industrial Engineering: Concepts, Methodologies, Tools, and Applications (pp. 1719-1733).*
www.irma-international.org/chapter/creative-regions-europe/69362

The Future of Product Design Education Industry 4.0
Jennifer Loyand James I. Novak (2021). *Research Anthology on Cross-Industry Challenges of Industry 4.0 (pp. 1666-1685).*
www.irma-international.org/chapter/the-future-of-product-design-education-industry-40/276895

Optimization Methods in Continuous Improvement Models: A Relational Review
Brian J. Galli (2019). *International Journal of Applied Industrial Engineering (pp. 46-59).*
www.irma-international.org/article/optimization-methods-in-continuous-improvement-models/222795

Multi-Criteria Evaluation Approach of Mobile Text Entry Methods
Lanndon Ocampo (2014). *International Journal of Applied Industrial Engineering (pp. 1-21).*
www.irma-international.org/article/multi-criteria-evaluation-approach-of-mobile-text-entry-methods/138306