

Chapter 16

Excess Entropy in Computer Systems

Charles Loboz
Microsoft Corporation, USA

ABSTRACT

Modern data centers house tens of thousands of servers in complex layouts. That requires sophisticated reporting – turning available terabytes of data into information. The classical approach was introduced decades ago to handle a small number of lightly connected computers. Today, we also need to identify problematic groups of servers, strange patterns in load, and changes in composition with minimal human involvement. The authors show how, as a single concept, entropy can describe multiple aspects of system use. Entropy is well grounded in physics, used in economics, and the authors extend it to large computer systems.

INTRODUCTION

Complexity and scale of computer systems keeps growing. A modern server has over 2000 performance counters which are relevant to the description of its state and usage – that applies to both Windows and UNIX servers. For one server we can select a smaller subset of counters requiring monitoring, but if we have several servers running different applications the size of the monitoring set grows quickly. Modern data centers house tens of thousands of servers in complex layouts.

Global provision of services requires tens of data-centers - and many such data centers are required for global provisioning of services. That generates a large volume of data – but, more importantly, this data is both complex and not easily tractable by traditional methods.

The costs of the infrastructure so large run into hundreds of millions of dollars per data center and efficient use of that infrastructure requires sophisticated reporting management. That, in turn, requires turning available (terabytes) of data describing system use into information.

DOI: 10.4018/978-1-4666-4699-5.ch016

Computer system performance analysis and capacity planning started decades ago with a single mainframe. Then we have moved to multiple mainframes and groups of servers. That was followed by multiple virtual machines running on a single server. The current stage – cloud computing – is, in effect, an operating system controlling execution of processing on clusters of servers and cluster groups.

We need to consider both traditional system descriptors as well as the new ones arising from the handling server groups and virtualization. Examples of the need for new descriptors include effects of competition for disk bandwidth between multiple virtual machines running on the same server and sharing physical disks – or similar competition for network rack switches between virtual machines deployed to the same rack of servers.

Performance analysts and capacity planners have to deal with information explosion in two different dimensions. The first one is related to the scale of modern web services, when datacenters containing tens of thousands of servers are providing thousands of services – thus we have data from a single server multiplied. The second dimension is growing layering and complexity of the underlying data.

Most complexity in this dimension is coming from the number of servers. The second dimension is the virtualization and cloud artifacts – consideration of deployment strategy for virtual machines, consideration of migration options for virtual machines to other servers or clusters and management of whole clusters of servers. To manage such information explosion we need descriptors of overall system usage that are on higher conceptual level than direct performance counters, like processor utilization, number of disk operations, memory bytes used, packets transferred through a network and other performance counters of this type.

Classical methods of describing and analyzing the use of computer systems were introduced decades ago (Lazowska, 1984), (Jain, 1991) and

designed to handle a small number of lightly connected computers. Introduction of new methods is forced by the need to handle problems arising from the growing size and complexity of new systems – operations in such system require higher-level descriptors.

An example of such a higher-level descriptor is Performance Impact Factor (PIF) introduced in (Loboz, 2009). For servers the average processor utilization is frequently misleading, because a low daily average can hide occasional spikes during the day – and such spikes may create reduced response time with disastrous consequences to service level agreements. PIF was designed to summarize in one number existence of such spikes. That replaces the need for looking at daily load charts – clearly impractical even with thousands of servers. In effect PIF transforms the data from performance counter space to performance impact space. That simplifies analysis of a large number of servers, because PIF is a one-number summary and captures the information not easily discernible from the original counters. PIF does not replace the traditional utilization description – it augments it and aggregates it so handling of a very large number of servers becomes practical.

Another example of a scalable aggregating descriptor is Capacity Utilization Factor introduced in (Loboz, 2010). It allows comparison of usage levels between servers with different hardware and between groups of such servers.

Describing usage of a system consisting of tens of thousands of computers with classical descriptors, like utilization coefficients or response times, is still necessary. However, with at the current scale and complexity we need to do much more. We must identify problematic groups of servers, strange patterns in load, discover changes in composition of customers, classify historical use – and do that with minimal human involvement. That requires new descriptors and approaches.

This paper outlines an approach to describe many aspects of system use with one concept: entropy. The concept is well grounded in science

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/excess-entropy-in-computer-systems/85465

Related Content

A Roadmap on Updates

Fernando Zacarías Flores, Dionicio Zacarías Flores, Rosalba Cuapa Canto and Luis Miguel Guzmán Muñoz (2009). *Database Technologies: Concepts, Methodologies, Tools, and Applications* (pp. 2261-2267).

www.irma-international.org/chapter/roadmap-updates/8034

Materialized Views in Multidimensional Databases

Stefano Paraboschi, Giuseppe Sindoni, Elena Baralis and Ernst Teniente (2003). *Multidimensional Databases: Problems and Solutions* (pp. 222-251).

www.irma-international.org/chapter/materialized-views-multidimensional-databases/26970

The Expert's Opinion

Mohammad Dadashzadeh (1991). *Journal of Database Administration* (pp. 37-41).

www.irma-international.org/article/expert-opinion/51089

Real World Knowledge for Databases

Veda C. Storey (1992). *Journal of Database Administration* (pp. 1-21).

www.irma-international.org/article/real-world-knowledge-databases/51098

Knowledge Discovery and Geographical Databases

Sami Faiz (2005). *Encyclopedia of Database Technologies and Applications* (pp. 308-312).

www.irma-international.org/chapter/knowledge-discovery-geographical-databases/11165