An ultra-large-scale simulation framework ☆

Dhananjai M. Rao and Philip A. Wilsey*

Experimental Computing Laboratory, Department of Electrical & Computer Engineering, Computer Science University of Cincinnati, Cincinnati, OH 45221-0030, USA

Abstract

Many modern systems involve complex interactions between a large number of diverse entities that constitute these systems. Unfortunately, these large, complex systems frequently defy analysis by conventional analytical methods and their study is generally performed using simulation models. Further aggravating the situation, detailed simulations of large systems will frequently require days, weeks, or even months of computer time and lead to scaled down studies. These scaled down studies may be achieved by the creation of smaller, representative, models and/or by analysis with short duration simulation exercises. Unfortunately, scaled down simulation studies will frequently fail to exhibit behaviors of the full-scale system under study. Consequently, better simulation infrastructure is needed to support the analysis of ultra-large (models containing over 1 million components)-scale models.

Simulation support for ultra-large-scale simulation models must be achieved using low-cost commodity computer systems. The expense of custom or high-end parallel systems prevent their widespread use. Consequently, we have developed an Ultra-large-Scale Simulation Framework (USSF). This paper presents the issues involved in the design and development of USSF. Parallel simulation techniques are used to enable optimal time versus resource tradeoffs in USSF. The techniques employed in the framework to reduce and regulate the memory requirements of the simulations are described. The API needed for model development is illustrated. The results obtained from the experiments conducted using various system models with two parallel simulation kernels (comparing a conventional approach with USSF) are also presented.

1. Introduction

Modern systems such as microprocessors and communication networks have steadily grown in size and sophistication to meet the ever increasing needs and demands. For example, today's microprocessors are built using a few million

---

*Corresponding author.

E-mail addresses: dmadhava@ececs.uc.edu (D.M. Rao), philip.wilsey@ieee.org (P.A. Wilsey).

transistors [14] and the Internet, a global data network, now connects more than 16 million nodes [15]. These systems involve complex interactions between a few thousand to several million entities. The study and analysis of these systems is necessary in order to effectively design, manufacture, and maintain them [15,25]. Unfortunately, analytical methods of analysis are insufficient to study these systems and experimental techniques such as computer-based simulations are usually employed instead [21,25]. Furthermore, parallel simulation techniques are employed to enable simulation of large systems in acceptable time frames [15,21,25]. Simulation enables explorations of complicated scenarios that would be either difficult or impossible to analyze [15]. Due to its effectiveness, simulation has gained considerable importance and is widely used today.

Validity of the models plays central role in analyzing systems using simulation [23]. The models should reflect the size and complexity of the system in order to ensure that crucial scalability issues do not dominate during validation of simulation results. Many techniques, algorithms, and protocols that work acceptably for small models consisting of tens or hundreds of entities may become impractical when the size of the system grows [15]. Events that are rare or that do not even occur in small toy models may be common in the actual system under study. Detailed simulation of the complete system is necessary to study large-scale characteristics, long-term phenomena, and to analyze the system as a whole. Paxson et al., provide an excellent context from the networking domain to highlight this issue. They write, "Indeed, the HTTP protocol used by the World Wide Web is a perfect example of a success disaster. Had its designers envisioned it in use by virtually the entire Internet—and had they explored the corresponding consequences with experiments, analysis or simulation—they would have significantly altered its design, which in turn would have led to a more smoothly operating Internet today" [15]. Since today's systems involve a large number of entities ranging in the order of a few millions, modeling and simulating such ultra-large systems is necessary.

Simulation of large systems is complicated due to their sheer size. The memory and computational resources needed to simulate such large systems in acceptable time frames are often beyond the limits of a single stand alone workstation [18]. Developing large and complex models while paying special care to optimally utilize system resources (in particular, memory) is a tedious task demanding considerable expertise from the modeler. Parallel simulation techniques need to be efficiently exploited to meet the computational requirements. However, investing in large and expensive hardware components for a "one time" analysis of the system models is seldom economically viable. Hence, simulating large systems using modest hardware resources is an attractive and often the only alternative.

This paper presents the design and evaluation of an Ultra-large-Scale Simulation Framework (USSF) that was developed to enable and ease effective simulation of large systems. In particular, USSF was motivated by the need to support analyzing systems involving millions of entities using only modest computational resources. The framework utilizes parallel simulation techniques to harness the resources of conventional workstations to provide optimal time versus resource tradeoffs. Various software techniques have been employed to reduce and regulate the memory requirements of the simulations. USSF provides a flexible and robust object-oriented API for model development. The API also insulates the model developer from the intricacies of enabling large simulations.

The remainder of this paper is organized as follows. A brief description about the parallel simulation kernels that are used as the underlying synchronization kernels of USSF are presented in Section 2. In Section 3, brief descriptions about some of the earlier research activities related to large-scale simulations are presented. Section 4