ARTICLE IN PRESS

Future Generation Computer Systems ▮ (▮▮▮▮) ▮▮▮–▮▮▮

Contents lists available at ScienceDirect

Future Generation Computer Systems

journal homepage: www.elsevier.com/locate/fgcs

Multiuser-centered resource scheduling for collaborative display wall environments

Sungwon Nam *, Khairi Reda, Luc Renambot, Andrew Johnson, Jason Leigh

Electronic Visualization Laboratory, 842 W. Taylor St., Chicago, IL 60607, United States

HIGHLIGHTS

- We present a model that prioritizes applications based on how they are presented.
- We propose a resource scheduling scheme that achieves presentation fairness.
- User study evaluates the proposed scheduler in a multiuser collaborative session.

ARTICLE INFO

Article history:
Received 30 November 2012
Received in revised form
10 March 2014
Accepted 15 August 2014
Available online xxxx

Keywords:
Scheduling
Human factors
Algorithms
Interactive systems
Distributed graphics

ABSTRACT

The popularity of large-scale, high-resolution display walls, as visualization endpoints in eScience infrastructure, is rapidly growing. These displays can be connected to distributed computing resources over high-speed network, providing effective means for researchers to visualize, interact with, and understand large volumes of datasets. Typically large display walls are built by tiling multiple physical displays together and running a tiled display wall required a cluster of computers. With the advent of advanced graphics hardware, a single computer can now drive over a dozen displays, thereby greatly reducing the cost of ownership and maintenance of a tiled display wall system. This in turn enables a broader user base to take advantage of such technologies. Since tiled display walls are also well suited to collaborative work, users tend to launch and operate multiple applications simultaneously. To ensure that applications maintain a high degree of responsiveness to the users even under heavy use loads, the display wall must now ensure that the limited system resources are prioritized to maximize interactivity rather than thread-level fair sharing or overall job-completion throughput. In this paper, we present a new resource scheduling scheme that is specifically designed to prioritize responsiveness in collaborative large display wall environments where multiple users can interact with multiple applications simultaneously. We evaluate our scheduling scheme with a user study involving groups of users interacting simultaneously on a tiled display wall with multiple applications. Results show that our scheduling framework provided a higher frame-rate for applications, which led to a significantly higher user performance (approx. 25%) in a target acquisition test when compared against traditional operating system scheduling scheme.

1. Introduction

Data-intensive eScience applications are driven by global cyberinfrastructure where visualization platforms, computing and storage resources and instruments are distributed and interconnected with high-speed network. The vast amount of data produced by these eScience applications creates a major challenge for researchers who must manage and interpret the increased scale of their work.

Visualization provides effective means in scientific discovery process. One can verify the correctness of a complex simulation model, provide more insight into the model, and present results in a way that it can be more easily understood. Large-scale, high-resolution display walls are used in scientific disciplines because they are an effective way to provide both context and details when visualizing high-resolution data. Furthermore the expansive size and exquisite resolution of the display has been conclusively shown to positively impact the scientific discovery process by allowing researchers to juxtapose multiple high-resolution visualizations simultaneously [1–8].

* Corresponding author. Tel.: +1 312 996 3002.
E-mail addresses: snam5@uic.edu, schiktol@gmail.com (S. Nam), mreda2@uic.edu (K. Reda), renambot@uic.edu (L. Renambot), ajohnson@uic.edu (A. Johnson), spiff@uic.edu (J. Leigh).

http://dx.doi.org/10.1016/j.future.2014.08.012
0167-739X/

Fig. 1. A single machine-driven tiled-display wall running SAGE at the Electronic Visualization Laboratory in the University of Illinois at Chicago. The 20′ by 6′ display wall is made up of 18 LCD panels with a total resolution approximately 17 megapixels.

A common way to build large-scale display walls is to tile multiple individual displays and connect them to a cluster of computers. Thus cluster middleware is needed to enable users to work with the wall as a single contiguous display surface. Traditional tiled display middleware such as CGLX [9], Chromium [10] and Equalizer [11], that are designed as large-scale visualization platforms, can be regarded as distributed graphic frameworks. They focus on parallel rendering of large-scale datasets using a computer cluster, and are typically aimed at cases where a single user interacts with a single application spanning the entire display wall.

On the other hand, tiled display middleware such as the scalable adaptive graphics environment (SAGE) [12] lets users launch distributed visualization applications on remote clusters whose outputs are then streamed directly to display walls. This makes SAGE a low-cost, "thin-client" visualization endpoint where such visualizations are rendered by remote computing resources and streamed over an optical network to a display wall. SAGE also provides highly collaborative visualization environments by enabling multiple users to simultaneously view and interact with these applications on large-scale display walls [13]. An overview as well as real-world use cases of the thin-client display wall paradigm is discussed in [14,15]. In [16], we show how SAGE coupled with ParaView [17] can be used as a thin-client display for scientific visualization applications.

The emergence of multi-headed graphic technologies (such as NVIDIA's Scalable Visualization Solutions and AMD's Eyefinity), has greatly amplified the graphical capabilities of a single computer node. This empowers a single computer to drive a large-scale display wall, in many cases eliminating the need for a computer cluster, which significantly reduces the cost of ownership and maintenance of these environments. Furthermore, applications can now run natively on a single-machine without the need to parallelize them, thus simplifying application development for large-scale display walls. Fig. 1 shows a 20×6 foot large-scale tiled display wall driven by a single computer machine at the Electronic Visualization Laboratory in the University of Illinois at Chicago.

Driving a multiuser collaborative large-scale tiled display wall with a single computer however, presents significant challenges in resource management. A display wall middleware relying on general-purpose operating system resource scheduling may fail to provide a good user experience in collaborative large-scale display environments where multiple users interact simultaneously with multiple applications. A general-purpose operating system schedules resources based on system-wide performance measures such as job completion throughput and fine-grained fairness. In large-scale collaborative display wall environments, multiple users may simultaneously view and interact with Cloud media data such as pictures, documents, and movies, VNC-shared desktop screens, and interactive scientific visualization. Users can also move, resize, and arrange windows on the display wall in a variety of layouts. The number of applications running on the system, their layouts on the display, and the user-interaction pattern in these systems can differ drastically from traditional desktop environments where a single user typically interacts with a limited number of applications. This difference makes traditional resource scheduling schemes unfit for tiled display wall environments.

Fig. 2 shows examples of layouts on a large-scale display wall with varying degrees of window overlap. A traditional operating system will try to ensure fair sharing of system resources in all cases depicted in Fig. 2, while fair sharing might only be useful in the case depicted in (a). When the window layout is arbitrary as in (b), giving more system resources to windows with which users are interacting can achieve a better user experience than a fair sharing. Similarly, a better user experience can be achieved in (c) if more system resources are allocated to the application whose window is in the foreground. For the case shown in (d), a fair-sharing scheme

Please cite this article in press as: S. Nam, et al., Multiuser-centered resource scheduling for collaborative display wall environments, Future Generation Computer Systems (2014), http://dx.doi.org/10.1016/j.future.2014.08.012