Future Generation Computer Systems 54 (2016) 336–343

Contents lists available at ScienceDirect

Future Generation Computer Systems

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

Low bandwidth desktop and video streaming for collaborative tiled display environments

Jason Kimball *, Tom Wypych, Falko Kuester

Center of Interdisciplinary Science for Art, Architecture and Archaeology (CISA3), Qualcomm Institute, Calit2, University of California, San Diego, United States

HIGHLIGHTS

- H.264 video streaming of desktop content to high resolution tiled display systems.
- Low bandwidth and low latency streaming outperform existing raw RGB techniques.
- Enables streaming full HD resolution desktop content from wireless laptops.
- Removes the dependence on 10 Gbps networks in collaborative tiled display systems.
- Demonstration system delivers 1080P30 desktop content under 10 Mbps, 100 ms latency.

ARTICLE INFO

Article history:
Received 20 July 2011
Received in revised form
16 July 2015
Accepted 18 July 2015
Available online 30 July 2015

Keywords:
Real-time video streaming
Tiled display environments
Collaborative visualization
H.264 video compression

ABSTRACT

High-resolution display environments built on networked, multi-tile displays have emerged as an enabling tool for collaborative, distributed visualization work. They provide a means to present, compare, and correlate data in a broad range of formats and coming from a multitude of different sources. Visualization of these distributed data resources may be achieved from a variety of clustered processing and display resources for local rendering and may be streamed on demand and in real-time from remotely rendered content. The latter is particularly important when multiple users want to concurrently share content from their personal devices to further augment the shared workspace. This paper presents a high-quality video streaming technique allowing remotely generated content to be acquired and streamed to multi-tile display environments from a range of sources and over a heterogeneous wide area network.

The presented technique uses video compression to reduce the entropy and therefore required bandwidth of the video stream. Compressed video delivery poses a series of challenges for display on tiled video walls which are addressed in this paper. These include delivery to the display wall from a variety of devices and localities with synchronized playback, seamless mobility as users move and resize the video streams across the tiled display wall, and low latency video encoding, decoding, and display necessary for interactive applications. The presented technique is able to deliver 1080p resolution, multimedia rich content with bandwidth requirements below 10 Mbps and low enough latency for constant interactivity. A case study is provided, comparing uncompressed and compressed streaming techniques, with performance evaluations for bandwidth use, total latency, maximum frame rate, and visual quality.

1. Introduction

Tiled display environments offer high resolution display resources at a scale unparalleled by other display technologies. These higher resolution display surfaces allow for larger datasets to be visualized on larger workspaces and support new modes of collaboration among users. Several techniques currently exist for filling the high pixel count with certain types of data, but challenges still exist for displaying many types of content on tiled display environments, such as a wide range of desktop applications and tools that researchers are already familiar with, and sharing content from multiple sources which are external to the display environment.

While projects like the OptiPortal research initiative [1] that are developing cost-effective tiled display wall technology from commodity components aim to increase access to tiled display walls, wide spread adoption of tiled display environments will

* Corresponding author.
E-mail addresses: jkimball@eng.ucsd.edu (J. Kimball), twypych@ucsd.edu (T. Wypych), fkuester@ucsd.edu (F. Kuester).

http://dx.doi.org/10.1016/j.future.2015.07.009

not occur until users can access and visualize their data in a way that they are familiar with. Despite the benefits higher resolution display walls provide, an interim solution is necessary until these tiled display walls support all of the visualization functionality researchers need.

Until now a popular solution has been to stream remote content directly from rendering nodes in an uncompressed format that requires high bandwidth interconnects. The high bandwidth required for this streaming limits the types of devices that can provide content, as well as the total resolution and update rate of the content streamed.

To address the issue of high speed network dependence we present a new framework for streaming remote content to tiled display environments using low latency H.264 video compression. The use of video compression on a video stream significantly reduces the required bandwidth and network resources which provides several improvements over uncompressed streaming. Real-time video streaming capabilities become accessible to a new class of bandwidth constrained devices such as wireless laptops. Also, the number of concurrent video streams that can be sent over existing high speed network infrastructure is increased, allowing collaboration with multiple concurrent content streams on a network with a 1 Gbps bandwidth or significantly less.

This paper describes the implementation of a complete end-to-end system for streaming desktop content to tiled display walls as outlined in Fig. 1. It also addresses the challenges of streaming compressed video to a tiled display wall including acquisition of a range of source content, extremely low-latency video encoding, efficient network transport over heterogeneous wide area networks, video decompression, and visualization on a tiled display system. Also addressed is the issue of mobility, allowing the users to dynamically reposition and resize each video stream anywhere on the tiled display wall without interruption.

The presented framework provides support for the acquisition and display of a variety of desktop and video content, allowing users the freedom to visualize and collaborate on tiled display walls with conventional desktop applications, HD video cameras, video game consoles, and almost any device which outputs to an HDMI or DVI interface.

In the following sections we describe the components of the system that make this possible and present a performance analysis between streaming of uncompressed RGB pixel streams, single frame compressed streams, and H.264 compressed streams. We are able to demonstrate that using the system outlined in this paper, H.264 compressed video streams provide a higher frame throughput, lower end-to-end latency, and require significantly less bandwidth than uncompressed RGB streams while providing a higher quality and lower bandwidth usage than competing real-time single frame compression approaches.

Fig. 1. An example of 4 live video streams on a 32 screen tiled display wall. Researchers and a remote collaborator compare a collection of images from previous San Diego County fires with video streams of a real-time wind flow visualization, current fire locations map on a tablet, and live panoramic images from various San Diego County locations on a laptop.

2. Related work

This paper bridges contributions from two disjoint areas of research: displaying content on tiled display walls, and streaming desktop content for collaboration.

2.1. Scalable tiled display wall applications

Several existing middleware systems provide access to different content in tiled display environments. Chromium [2] and DMX [3] are the most non-invasive middleware in terms of allowing generic applications to be run. They both operate by intercepting graphical API calls of an application running on a dedicated head node and rerouting them to a distributed display application. These approaches restrict scalability because content has to be distributed from one computer running the application to the display wall, limiting the amount of content to that which can be processed by the head node. Furthermore, Chromium only works with OpenGL applications, and can only run one application at a time, and DMX only works with X11 environments and has limited support for hardware accelerated graphics such as OpenGL. While both of these approaches allow easy access to certain subsets of existing applications, both are limited in scalability and neither are a solution for collaborative visualization as neither approach can bring together content from multiple source computers.

SAGE is a pixel streaming middleware which takes raw RGB input from a source application and streams it to a tiled display wall [4]. It can display simultaneous content from multiple sources and those sources can be subdivided across a rendering cluster in order to parallelize applications which are data, CPU, or graphics intensive. In this way, SAGE improves upon scalability by allowing rendering to be parallelized and enables collaborative input by displaying streams from multiple sources. Because SAGE streams raw RGB pixel data, bandwidth usage can be very high. Applications streaming high resolution content to multiple displays require 10 Gbps interconnects and networking hardware with even higher bisection bandwidth as the aggregate bandwidth to all of the displays can well exceed 10 Gbps. Furthermore, applications must be recompiled to take use of the SAGE Application Interface Library (SAIL).

Careful segmentation of the SAGE stream is required in order to conserve the bandwidth of data sent to each display node and computational effort is required to segment the video stream each frame. SAGE facilitates the segmentation of the rendered stream geometry to match the output geometry on the tiled display wall. This requires communication between the rendered source and the display nodes whenever users want to move or resize content on the display wall, resulting in mobility latency. In this paper's approach, the video bandwidth is much lower and can be distributed to all nodes, so there is no delay when moving or resizing content.

For collaboration between multiple SAGE display walls, Renambot et al. [5] implement a network utility, called SageBridge, which dynamically re-segments the rendered streams when streaming to multiple display walls. Jeong et al. [6] improve on the collaborative display of SAGE streams by introducing macro-block segmentation as a replacement for per-pixel image segmentation across multiple displays. This simplifies the task of segmenting and distributing a large stream, reducing the load on the SageBridge machines when collaborating between a large number of display walls. CGLX [7] is