Available online at www.sciencedirect.com
SCIENCE DIRECT®
Signal Processing: *Image Communication* 21 (2006) 396–415
SIGNAL PROCESSING: IMAGE COMMUNICATION
www.elsevier.com/locate/image

On-demand transmission of 3D models over lossy networks

Dihong Tian, Ghassan AlRegib*

School of Electrical and Computer Engineering, Georgia Institute of Technology, 210 Technology Circle, Savannah, GA 31407-3039, USA

Received 7 September 2005; accepted 25 January 2006

Abstract

Three-dimensional (3D) meshes are used intensively in distributed graphics applications where model data are transmitted on demand to users' terminals and rendered for interactive manipulation. For real-time rendering and high-resolution visualization, the transmission system should adapt to both data properties and transport link characteristics while providing scalability to accommodate terminals with disparate rendering capabilities. This paper presents a transmission system using hybrid unequal-error-protection and selective-retransmission for 3D meshes which are encoded with multi-resolutions. Based on the distortion-rate performance of the 3D data, the end-to-end channel statistics and the network parameters, transmission policies that maximize the service quality for a client-specific constraint is determined with linear computation complexity. A TCP-friendly protocol is utilized to further provide performance stability over time as well as bandwidth fairness for parallel flows in the network. Simulation results show the efficacy of the proposed transmission system in reducing transmission latency and providing smooth performance for interactive applications. For example, for a fixed rendering quality, the proposed system achieves 20–30% reduction in transmission latency compared to the system based on 3TP, which is a recently presented 3D application protocol using hybrid TCP and UDP.

Keywords: Distributed 3D graphics; Interaction; Multi-resolution compression; Unequal error protection; Selective retransmission; TCP-friendly rate control; Multimedia streaming

1. Introduction

Internet-based multimedia applications are expanding from streaming video/audio to distributed 3D graphics, driven by growing demands of various applications such as electronic commerce, collaborative CAD, medical and scientific visualization, and virtual environments. Interaction is one of the key aspects of a distributed graphics application. Most of the distributed graphics applications require certain level of interaction with the objects involved. Some applications such as immersive environments and shared reality thoroughly involve users in the interaction. Response time, which is the latency between the user input and the response (e.g., scenes displayed on the user's terminal) from the system, is one of the major considerations in designing high-performance distributed graphics systems. In contrast to specific 3D systems that assume all models are locally available and are essentially designed as stand-alone systems, distributed graphics applications often require on-demand exchange of 3D data in a networked environment, and impose requirements of real-time response and

*Corresponding author. Tel.: +1 912 966 7937; fax: +1 912 966 7928.
E-mail addresses: dhtian@ece.gatech.edu (D. Tian), gregib@ece.gatech.edu (G. AlRegib).

0923-5965/$ - see front matter
doi:10.1016/j.image.2006.01.003

smooth performance over time on the interaction. In addition, 3D data differ from general media content for it requires rendering capability on the user terminal. Sending the same dimension of data to user terminals with disparate rendering capabilities may result in significant difference in response time between the different clients. Scalability of the 3D data, as well as the transmission, is therefore desirable. All the aforementioned requirements, not only have promoted the use of high-performance computing systems and distributed platforms, but also call for careful considerations on ways of reducing transmission latency, providing scalability, and maintaining high-resolution visualization when network delays and random data losses are involved.

Multi-resolution compression of 3D meshes [1,16,19–21,23,24] is a partial solution to provide scalability for 3D data. Using multi-resolution encoders, the server can select the appropriate resolution for a particular client according to its quality requirement, or initially sends a coarse representation of the 3D model to the client for quick reconstruction and rendering, and then transmits refinement layers which allow the client to gradually increase model fidelity toward higher resolutions. Although, such methods are successful in exploring the space and time efficiency of 3D data, higher efficiency can be accomplished by addressing the effects of network behaviors. In particular, to display 3D scenes on the user's terminal with satisfactory quality and in real time, the impact of packet losses and transmission delays on the decoding process need to be explored. Typically, reliable or error-resilient transmission can be achieved by pre-processing techniques such as data partitioning [7,17,27], post-processing techniques such as error-concealment, and network-oriented techniques such as forward error correction [2,6,4] and retransmission techniques [3,5,8,9,14]. All these techniques address efficient transmission of 3D data *separately*. Yet an appropriate combination of such techniques is desired to achieve better performance. Interaction and tradeoffs among the selected techniques need to be investigated, taking into account the distortion-rate performance of the 3D data and the network characteristics.

In this paper, a hybrid mechanism of unequal-error-protection and selective-retransmission is proposed for multi-resolution meshes. Hierarchical data batches of the multi-resolution mesh are protected preferentially according to their distortion-rate performance, network parameters, and channel statistics estimated by the transport layer. To minimize response time in interaction, the proposed mechanism is designed to have linear computational complexity. In addition, by integrating TCP-friendly congestion control into the system, the proposed mechanism achieves smooth performance over time as well as bandwidth fairness for co-existing applications in the network. Simulation results show the efficacy of the proposed mechanism. For instance, compared with a recently presented 3D application protocol referred to as 3TP [5], the proposed system achieves 20–30% reduction in transmission latency while delivering the same level of rendering quality.

The main contributions of the presented work are summarized as follows:

- We analyze intensively the property of multi-resolution 3D meshes and present a TCP-friendly transmission system for multi-resolution meshes including a novel and meaningful quality measure.
- Given a distortion constraint, we derive a simplified expression (with assumptions) of the optimal FEC code that minimizes the expected transmission latency when combined with retransmissions.
- Observing a semi-infinite space for finding the theoretical optimal solution, we propose an extended steepest decent search algorithm which quickly reaches the local optima in the solution space.
- Based on the above results for quality-critical scenarios, we further develop a time-critical streaming algorithm which significantly decreases the receiving distortion upon a strict delay constraint.

The rest of the paper is organized as follows. The relevant prior art that addresses channel effects in multi-resolution mesh transmission is summarized in the next section. Major aspects of the proposed mesh transmission system are described in Section 3, and Section 4 presents a detailed study of the hybrid unequal-error-protection and selective-retransmission mechanism. Test results in simulated network environments are given in Section 4. Finally, Section 5 concludes the paper and summarizes future work.