

A Multimedia Information Server with Mixed Workload Scheduling

Guido Nerjes

Department of Computer Science, University of the Saarland, Germany

E-Mail: nerjes@cs.uni-sb.de, WWW: www-dbs.cs.uni-sb.de

Abstract

In contrast to specialized video servers, advanced multimedia applications for tele-shopping, tele-teaching and news-on-demand exhibit a mixed workload with massive access to conventional, “discrete” data such as text documents, images and indexes as well as requests for “continuous data” such as video. This paper briefly describes the prototype of a multimedia information server that stores discrete and continuous data on a shared disk pool and is able to handle a mixed workload in a very efficient way.

1. Motivation

Advanced multimedia applications that make use of both discrete data (text, html, graphics) and continuous data (audio, video) include tele-shopping, tele-teaching or news-on-demand applications. In contrast to specialized video servers [1] and unlike conventional web servers, multimedia information servers have to deliver both data types to the clients. Furthermore, with unrestricted 24-hour world-wide access over the Web, multimedia information servers have to cope with a dynamically evolving workload where the fractions of continuous-data versus discrete data requests vary over time and cannot be completely predicted in advance. Thus, for a good cost/performance ratio it is mandatory that such a server operates with a shared resource pool rather than statically partitioning the available disks and memory into two pools for continuous and discrete data, respectively.

Quality of service requirements for “continuous” data like video and audio pose challenging performance demands on a multimedia information server. In particular, the delivery of such data from the server to its clients dictates disk-service deadlines for real-time playback at the clients. Missing a deadline may result in a temporary, but possibly user-noticeable degradation of the playback that we refer to as a “glitch”. Guaranteeing a specified quality of service then means to avoid glitches or to bound the glitch rate within a continuous data stream, possibly in a stochastic manner (i.e., with very high probability) [3]. In addition to the service quality guarantees for continuous data requests (*C-requests*), quality-conscious applications require that the response time of the discrete data requests (*D-requests*) stay below some user-tolerance threshold, say one or two seconds [6].

The prototype of the mixed workload multimedia information server developed at the University of the Saarland takes service quality requirements for both data types into account and implements a novel disk scheduling strategy that optimizes disk accesses. The prototype system

consists of three components that are briefly described in the following sections.

2. Architecture of the Prototype System

2.1 Mixed-Workload Multimedia Information Server

The mixed workload server stores continuous and discrete data objects on a shared disk pool. A continuous data object, e.g., a video, is partitioned into *fragments* of constant time length, say 1 second of display. These fragments are then spread over the disks in a round-robin manner such that each fragment resides on a single disk. Such a coarse-grained striping scheme [2] allows a maximum number of concurrent streams for a single object (regardless of skew in the popularity of objects), while also maximizing the effective exploitation of a single disk's bandwidth (i.e., minimizing seek and rotational overhead). Furthermore, the fact that all fragments have the same time length makes it easy to support data with variable-bit-rate encoding (e.g., MPEG-2) and simplifies the disk scheduling as follows. The periodic delivery of the fragments of the ongoing data streams is organized in *rounds* whose length corresponds to the time length of the fragments. During each round, each disk must retrieve those of its fragments that are needed for a client's playback in the subsequent round. Not being able to fetch all the necessary fragments by the end of a round is what causes a glitch. On the other hand, since the ordering of the fragment requests within a round can be freely chosen, the disk scheduling employs a scan policy (also known as “elevator” or “sweep” policy) that minimizes seek times.

In contrast, a discrete data object is placed completely on a single disk. The mixed workload server schedules the service of discrete data accesses within the rounds in a way such that continuous data requests are not affected. The scheduling strategy developed to this end is a mixed, dynamic, incremental scan strategy [4,5]. This strategy mixes *C-requests* and *D-requests* within a scheduling round and is able to avoid glitches for the *C-data* streams by dynamically limiting the number of *D-requests* that are served in a scheduling round. This makes the best possible use of the remaining disk time for good response time of *D-requests*. Figure 1 illustrates this strategy in a simple scenario. The timeline is split into three rounds of length T . Each box shows the beginning and the end of a request execution. *C-requests* are shown as striped boxes, *D-requests* as light shaded boxes. The number in a box indicates the cylinder number where the data resides on disk. Arrows above these boxes indicate the arrival of a *D-request*; arrows below the boxes show that a *D-request* has finished and leaves the system. In each round there are two *C-requests* that must be served. During the first round the *D-request* with the number

