

Online Scheduling in Modular Multimedia Systems with Stream Reuse*

Michael K. Bradshaw, Jim Kurose, Prashant Shenoy, Don Towsley

Dept. of Computer Science
University of Massachusetts
Amherst, MA 01003
{bradshaw, kurose, shenoy,towsley}@cs.umass.edu

May 12, 2005

UMass Computer Science Technical Report 2005-21

Abstract

When properly constructed, a modular multimedia system can satisfy a client's request in multiple ways by using different sequences of modules and by reusing existing streams within the system. Such flexibility in a multimedia server or proxy can provide a rich set of services to clients while efficiently utilizing system resources by choosing the best way to schedule (i.e. allocate resources for) new clients. However, it is difficult to optimally schedule clients in an online fashion as the problem is NP-complete. In this paper, we provide an efficient online algorithm to schedule client requests using a modular multimedia platform.

Categories and Subject Descriptors: C.4: Design Studies

General Terms: Algorithms, Design, Measurement

Keywords: Modular, Multimedia, Resource-Conserving

*This material is based on work supported by the National Science Foundation under Grant Nos. ANI-0085848 and EIA-0080119. Any opinions, findings conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

1 Introduction

The growing popularity of streaming media places an increasing strain on both network and server resources. Streaming media requires large amounts of network and disk bandwidth to successfully transmit streams from a server or proxy to the client. The most common solution is to simply purchase more resources. For instance, replicating servers and proxies in the Internet can increase both network and disk bandwidth [1, 5]. For every system there exists some workload that will completely consume one of the resources. In many of these cases, one set of resources can be substituted for another set of resources. Then, the system can adapt to the workload by exploiting resources that are not constrained.

Using excess resources to ease another, constrained resource is a common technique. Multimedia proxies cache streams in memory to conserve network bandwidth [19, 16]. Servers and proxies can transcode the stream, using processor time, to lower the network bandwidth needed [3]. Many of these techniques are passive, in that a designer studies the workload and designs the system to deliver maximum performance. Other techniques recognize changing workloads and adapt their strategies to the workload. Pooled RBC monitors the disk bandwidth to determine when caching streams is advantageous [2]. Interval caching determines when streams may be reused by a second client by caching a moving window of the stream in memory [7]. Energy-Aware systems trade off energy for quality of service [15]. In all existing work, the system's ability to adapt to resource contention is explicitly designed for a particular problem and is not generalized to generic resources types.

Modular multimedia systems are well-suited to provide resource conserving services due to their flexibility. *Stream modules* are used to implement each stream operation in the system. A module that reads a stream from the network and one that sends a stream into the network are two examples of stream modules. In order to use multiple modules in series, *pipes* are used to redirect the output stream of one module to serve as the input stream of another module. A sequence of modules is called a *pipeline*. Programmers can create new services by forming new pipelines. In order to react to fluctuations in resource consumption, a system [18] can use a lookup table to determine which pipeline should be used based on available resources. To avoid programming an exponential number of services (for each combination of modules), some systems [13, 14, 10] automate the construction of pipelines as requests arrive at the system. These systems search through stream module descriptions to find and instantiate a sequence of modules that can satisfy individual requests. While Ninja [9] does choose modules (actually nodes offering a service) based on the lowest CPU load, no modular multimedia system characterizes all of the resources used to determine the best way to satisfy a request.

AMPS [20] is a modular multimedia platform capable of sharing streams between clients. To enable sharing, pipes dynamically *branch* the output stream to serve as an input stream to more than one module. Branching allows the system to reuse the stream generated by previously instantiated modules to service multiple requests. In our earlier work [4], we introduced a Graph Manager (GM)

as a component of AMPS that composes modules online to satisfy requests. The GM selected a pipeline to satisfy each request by examining the set of resources needed by that pipeline and resources available in the system. We will show that finding an optimal solution is NP-complete in Section 4 (via reduction to the Interval Classroom Assignment problem [6]). To find good solutions at online speeds, we restricted how streams were shared between clients. That restriction allowed the GM to generate solutions quickly enough to support online scheduling.

In this paper, we provide a significant improvement in the GM's ability to schedule the system's resources to satisfy client requests. First, we formalize the problem of optimally scheduling an arriving client. Second, we present an exact search algorithm based on guided search to find the optimal schedule, and offer approximation techniques that enable the GM to efficiently schedule online. Finally, our results show that the GM can schedule a client very quickly, taking on average 40 ms per client, in the presence of 3600 concurrent clients. Furthermore, we present the resource savings realized by reusing existing streams to satisfy future clients.

The rest of the paper is structured as follows. In Section 2 we review the capabilities of the AMPS system. Then, we formulate the optimization problem based on the AMPS architecture in Section 3. We prove that the Stream Scheduling problem is NP-complete via reduction to the Interval Classroom Assignment problem in Section 4. In Section 5 we solve the optimization problem and provide the approximations needed to speed up the execution time. Next, we describe the evaluation environment in Section 6 and the experiments in Section 7. Finally, we conclude the paper in Section 8.

2 A Modular Multimedia Platform

AMPS [20] is a modular multimedia platform designed for rapid prototyping. AMPS has two components of interest to us: *stream modules* and *pipes*. Stream modules encode each stream service offered by the system such as: a transcoding service, a transmission service, or a service that generates added content. In order to use multiple modules in series, *pipes* are used to direct the output stream of one module to serve as the input stream of another module. By passing a stream through a sequence of stream modules, connected by pipes, the system can create a new service in the form of a *pipeline*. In this system, pipes can also dynamically *branch* the output stream to serve as an input stream to more than one module. Branching allows the system, typically a server or proxy, to reuse the stream generated by previously instantiated modules to service multiple requests.

Because of AMPS's modularity and branching, the system can satisfy requests by using different sequences of modules and by branching already scheduled pipelines. We present an example of AMPS's capabilities in Figure 1. Imagine a proxy with 3 modules: a Network Transmission (NT) module that sends a stream into the network, a Network Reception (NR) module that reads a

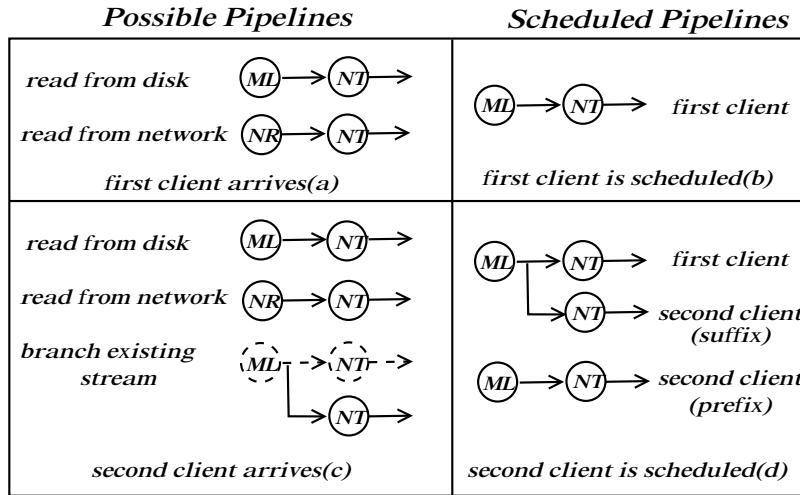


Figure 1: AMPS Capabilities: A modular multimedia platform with support for stream reuse

stream from the network and a Memory Loader (ML) module that reads a stream from disk. When the first request arrives at the system (Figure 1(a)), the request can be satisfied in one of two ways. The proxy can read the stream from disk and send it to the network (ML and NT), or the proxy can read the stream from the network and send it to the network (NR and NT). Suppose the proxy chooses to read the stream from disk (Figure 1(b)). When a second request arrives later (Figure 1(c)), there are 3 possible pipelines; the two previously mentioned and the pipeline that branches the previously scheduled pipeline to satisfy the suffix of the stream. Note that the prefix has already been generated and cannot be shared. In this case, the branched pipeline is chosen to satisfy the suffix of the stream, and a second pipeline is generated to supply the prefix (Figure 1(d)). We assume that the client can receive more than one stream at a time and can buffer frames that arrive early.

The AMPS system has a scheduler called the *Graph Manager* (GM) that determines the best way to satisfy requests. The GM's name comes from the graph-like appearance of the stream modules and pipes used to satisfy requests. The GM's task can be divided into two subtasks. First, the GM must find all possible pipelines that can be used, wholly or partially, to satisfy a new request, as seen in previous work [4]. Second, the GM must choose which pipeline or combination of pipelines to use in satisfying a request. When multiple pipelines can be used to satisfy a request, the system can use resources more efficiently, but scheduling multiple pipelines is much more complex. Our contribution focuses on the second subtask, namely that of determining which pipelines to schedule for each request.

3 Optimization Problem

In this section we present an optimization problem that arises in the AMPS system when a request arrives. The problem assumes that the system has served clients in the past and is concerned with minimizing the cost to the system of an arrival regardless of future arrivals. Note that a solution to the optimization problem may not lead to a “globally” optimal solution that accounts for all future client requests. However, in a real system, knowledge of future requests is not usually known.

In the rest of this section, we develop the full optimization problem in a series of steps. First, the *simplified problem* ignores some details/abilities of the AMPS system, but it provides the intuition for the heuristic driven search introduced in Section 5.1.2. Second, the *simplified problem with buffering* accounts for AMPS ability to buffer streams prior to delivery. Finally, the *full problem* captures all of the abilities of the AMPS system. In particular, it models the GM’s ability to schedule multiple portions of the stream at the same time.

3.1 Simplified Problem

A *request* v is a 3-tuple (b, e, t) that describes the deadline by which the GM must produce each frame. The GM must schedule all frames between $v.b$ and $v.e$. The first frame $v.b$ must be sent by the deadline $v.t$. All remaining frames $x, v.b \leq x \leq v.e$ have deadlines $v.t + (x - v.b)$.

In the simplified problem, the GM must decide how to satisfy the request when all frames are produced at their deadlines. Time is slotted. Each slot represents an interval of time, and each slot is indexed by an integer. Denote *current time* by t_0 . The system has k limited *resources* that are consumed when the GM generates a frame. A resource can be the bandwidth or the memory that is available to the system. Let $\vec{R}(t) = (R_1(t), \dots, R_k(t)); 0 \leq R_i(t) \leq 1$ represent the utilization of each of the k resources at time t . When we discuss resources, we always talk about what percentage of the total resources is used. For this reason, the system’s resources $\vec{R}(t) = (R_1(t), \dots, R_k(t))$ at time t are real values that lie between 0 and 1. The amount of resources available varies over time since the GM can experience different request rates, different types of requests, or choose to satisfy similar requests in different ways.

A *stream segment* is a potential, partial solution that has been generated by the GM. The choice of which modules to use (pipeline) and whether or not to make use of an existing stream are abstracted away (the algorithms needed to form stream segments are the basis of previous work [4]). A stream segment s is a 4-tuple (b, e, t, \vec{r}) that describes a stream that could be scheduled by the GM. The stream segment sequentially produces frames $s.b$ to $s.e$. The stream segment produces each frame $x, s.b \leq x \leq s.e$, at time $s.t + (x - s.b)$. $s.\vec{r} = (s.r_1, \dots, s.r_k); 0 \leq s.r_i$ is the vector of the k resources consumed to generate each frame in this stream segment. The values of $s.r_i, 1 \leq i \leq k$ represent

what fraction of the system's i^{th} resource is needed. When $s.r_i > 1$ the stream segment requires more resources than are present in the system.

\mathcal{S} is the set of all stream segments the system can produce to satisfy a request v . Given request v , all members $s \in \mathcal{S}$ must satisfy the following conditions:

$$s.t \geq t_0 \tag{1}$$

$$v.b \leq s.b < s.e \leq v.e \tag{2}$$

$$v.t + (x - v.b) = s.t + (x - s.b) \quad \text{for any frame } x \tag{3}$$

All members of \mathcal{S} only generate frames *at* the deadlines specified in the request v . Note that a particular stream segment might not contain all frames in a request due to branching of ongoing streams.

The cost function $c : [0, 1]^k \rightarrow \mathcal{R} \cup \{\infty\}$ maps a vector of resource utilizations, $\vec{r} = (r_1, r_2 \dots r_k; 0 \leq r_i)$, to a real number (or infinity if unsatisfiable). The cost function is:

$$c(\vec{r}) = \begin{cases} \sum_{i=1}^k r_i^2 & \text{if } \forall_{1 \leq i \leq k} r_i \leq 1 \\ \infty & \text{if } \exists_{1 \leq i \leq k} r_i > 1 \end{cases} \tag{4}$$

The cost function is the sum of the squares of the resource utilizations. The cost function is convex and increases faster than the resource utilization as the utilization approaches one. This cost increase discourages the GM from selecting stream segments that require heavily utilized resources and to instead favor stream segments that consume underutilized resources. If any of the resource utilizations is greater than one, the cost is ∞ and indicates that the system does not have the resources needed to satisfy the request.

The best solution not only satisfies the constraints of the request but also has the lowest sum of the marginal costs for each time slot. The optimization problem is to create a mapping $\mathcal{P} : x \rightarrow s$ for each frame $x, v.b \leq x \leq v.e$, to a stream segment $s \in \mathcal{S}$ that will generate that frame so as to minimize:

$$f(\mathcal{P}, \vec{R}(), t_0) = \sum_{t=t_0}^{\infty} c(\vec{R}(t) + \sum_{(x,s) \in \mathcal{P}} I(t = s.t + (x - s.b))s.\vec{r}) - c(\vec{R}(t)) \tag{5}$$

where $I()$ equals 1 if the argument is true and 0 otherwise. Note that not all stream segments can satisfy each frame in the request. The lowest cost schedule could use multiple stream segments to satisfy the request.

Since no two frames in the request are generated at the same time, the optimization problem can be restated to minimize:

$$f(\mathcal{P}, \vec{R}()) = \sum_{(x,s) \in \mathcal{P}} c(\vec{R}(s.t + (x - s.b)) + s.\vec{r}) \tag{6}$$

Equation (6) is used in the heuristic driven search presented in Section 5.1.2.

3.2 Simplified Problem with Buffering

In this section we extend the simplified problem to incorporate the ability to buffer stream data. In the simplified problem, streams that are branched for reuse are buffered until their delivery time. To model the resources consumed during the buffering period, we introduce changes to the definition of the stream segment and alter the function to be optimized.

A *stream segment* s is a 6-tuple $(b, e, t, \vec{r}, l, \vec{p})$ that describes a stream that could be generated by the GM. The meaning of b, e, t , and \vec{r} are the same as in the simplified problem. Two new components represent the properties of a stream segment with buffering. The lead time $s.l$ indicates the number of time slots each frame is buffered before it is produced. The preparation resources $s.\vec{p}$ indicate the resources needed to buffer frame x during the time segment $[s.t + (x - s.b) - l, s.t + (x - s.b) - 1]$.

The optimization problem is to create a set \mathcal{P} of pairings (x, s) for each frame $x, v.b \leq x \leq v.e$, with a stream segment $s \in \mathcal{S}$ so as to minimize:

$$f'(\mathcal{P}, \vec{R}(), t_0) = \sum_{t=t_0}^{\infty} c(\vec{R}(t) + \sum_{(x,s) \in \mathcal{P}} I(t = s.t + (x - s.b))s.\vec{r} + \mathbf{I}(s.t + (\mathbf{x} - s.b) - \mathbf{1} \leq \mathbf{t} \leq s.t + (\mathbf{x} - s.b) - \mathbf{1})s.\vec{p}) - c(\vec{R}(t)) \quad (7)$$

Where $I()$ equals 1 if the argument is true and 0 otherwise. The boldface (second) line of the equation is new and signifies the resources consumed through buffering. This is important, otherwise the system could schedule streams when there are not available resources (for instance, the memory required to buffer the streams prior to playback is not available).

3.3 Full Problem

The full problem allows the system to generate frames *before* their delivery deadlines. Frames can be generated before their playback time if the GM “branches” an existing stream and reuses it. Previously, the GM was forced to buffer the stream until its delivery deadline. In the full problem the GM can choose whether to buffer reused streams or immediately send them to the client. This seemingly minor difference makes the system more useful and yet much more difficult to solve.

In the full problem, we relax constraint (3) that forces all members of \mathcal{S} to generate all frames at their delivery deadlines. Given request v , all members $s \in \mathcal{S}$ must satisfy the following conditions:

$$\mathbf{s.t} - \mathbf{l} \geq \mathbf{t}_0 \tag{8}$$

$$v.b \leq s.b < s.e \leq v.e \tag{9}$$

$$\mathbf{v.t} + (\mathbf{x} - \mathbf{v.b}) \geq \mathbf{s.t} + (\mathbf{x} - \mathbf{s.b}) \text{ for any frame } \mathbf{x} \tag{10}$$

Changes to the simplified problem are in boldface.

The full problem must then minimize Equation (7) to find the optimal solution.

4 Complexity of the Optimization Problem

In this section we show that the complexity to optimize the full problem is NP-complete. To do this we first review a similar NP-complete problem called the Interval Classroom Assignment (ICA) Problem [6]. We then define the Stream Scheduling decision problem and show that any ICA problem can be reduced to a Stream Scheduling decision problem.

4.1 The Interval Classroom Assignment Problem

The Interval Classroom Assignment (ICA) problem schedules the rooms that a set of classes can meet in given time and room constraints for each class. This problem is seen every semester at high schools and places of higher education. Carter and Tovey showed that the ICA problem is NP-complete [6]. In the rest of this section we provide a more formal definition of the ICA problem.

The ICA problem is a tuple (N, M) that describes the set of classes N that need to be scheduled and the set of rooms M that classes can meet in. This problem uses slotted time. Each slot represents an interval of time and each slot is indexed by an integer. The slots could represent equal sized intervals (for instance 30 minute intervals) or each time slot could have different lengths. Each class is defined as a 3-tuple (b, e, A) where the class meets from time slot b to time slot e and A is the subset of rooms from M that are acceptable to schedule the class. We are asked if a schedule exists so that each class $n \in N$ can meet in a room $a \in n.A$ from times $n.b$ to $n.e$ and that no two classes are assigned to the same room during the same time slot.

4.2 The Stream Scheduling Decision Problem

The Stream Scheduling Decision Problem is derived from the optimization problem of the full problem in order to deliver a yes or no answer. In the Stream Scheduling Decision problem a


```

ICATransformation(N,M)
1   $k = |M|$ 
2   $\text{lastslot} = \max n.e; n \in N$ 
3  Forall  $t \leq \text{lastslot}$ 
4      $\vec{R}(t) = \vec{0}$ 
5   $x = 0$ 
6  Foreach  $n \in N$ 
7      $y = x + (n.e - n.b)$ 
8     Foreach  $a \in n.A$ 
9         create stream segment  $s = (x, y, n.b, \vec{a})$ 
10        place  $s$  into  $\mathcal{S}$ 
11     $x = y + 1$ 
12 create request  $v = (0, x, \text{lastslot})$ 

```

Figure 2: Polynomial Transformation of ICA problem to a Stream Scheduling problem

machine must determine if a set of stream segments \mathcal{S} can satisfy a request v mindful of the current resource utilizations $\vec{R}(t)$. In this section we show that the Stream Scheduling problem is in NP-Complete by first showing that the problem is NP and then by showing that the problem is NP-hard.

We prove that the Stream Scheduling problem is in NP by providing a certificate that can be verified in polynomial time. Our certificate is the set of stream segments that are used to satisfy the request. We can easily sort the stream segments and determine that all frames in the request have been generated once and only once in polynomial time. Then we only need to determine that no resource violations have occurred. This is accomplished by finding the cost of the solution $f'()$ as presented in Equation (7). (Note that Equation (7) does not need to evaluate values of t past the last frame deadline.) As the cost can be determined in polynomial time, the Stream Scheduling problem is in NP.

The final step is to show that the problem is NP-Hard. We will do this by showing that any Interval Classroom Assignment problem can be reduced to a Stream Scheduling problem. An algorithm of the reduction is presented in Figure 2. Let the tuple (N, M) be an ICA problem. Then, the stream scheduling problem will have $k = |M|$ resource types, one for each room. Initialize $\vec{R}(t)$ to the zero vector $\vec{0}$ for all reachable values of t (lines 2-4). This step is show for clarity and comparability with the full problem. The decision problem can be rewritten to set each value of $\vec{R}(t) = \vec{0}$ by default. Order the rooms in M . $R_i(t)$ corresponds to the use of the i^{th} room in M during time slot t . For the i^{th} room $m \in M$, let \vec{m} be the unit vector $= \langle 0, \dots, 1, \dots, 0 \rangle$ a zero vector with a one in the i^{th} element. For each class create a set of stream segments that have the same range of frames and starting times, but use different resources for each room (lines 6-11). Since the range of frames used

for each class is unique from the frames used for any other class, the Stream Scheduling algorithm will only select one of the stream segments generated for each class. Finally, a request is generated that covers the frames used for all of the classes, and has a deadline starting after all of the frames have been generated. The deadlines are not important. The running time of this transformation is $\mathcal{O}(|N||M|)$ as we only perform one operation for each room $m \in M$ that can be used by a class $n \in N$. (Note that lines 2-4 are only present for clarity). Therefore the Stream Scheduling problem is NP-hard and by the previous paragraph also NP-complete.

5 Algorithmic Solutions to the Optimization Problem

In this section we present the algorithms to solve the optimization problem in Section 3.3. First, we design an optimal algorithm to solve the optimization problem. While this algorithm will always find a solution, we show in Section 7.1 that it is not able to schedule solutions quickly enough to be used in real-time. Then to speed up the algorithm, we introduce two approximation techniques that reduce the search space and reduce the calculation times.

5.1 Exact Algorithms

We build our exact algorithm by first defining a search algorithm for the simplified problem. We then use this search algorithm as a heuristic to guide the search for the solution to the full problem. Note that the buffered problem is implicitly handled by the full search algorithm.

5.1.1 Simplified Search Algorithm

The simplified problem can be solved in a straightforward manner. Each mapping of a frame to a stream segment can be solved independently of the other frames see Equation (6). To map frame x to a stream segment, the GM searches through all stream segments and chooses the segment s with the lowest cost $c(\vec{R}(s.t + (x - s.b)) + s.\vec{r}); s.b \leq x \leq s.e$.

5.1.2 Full Search Algorithm

The full problem is more difficult to solve in that the cost to generate a frame depends on how many other frames are generated during the same time slot. In essence the cost of a solution cannot be calculated until each frame in the request has been mapped to a stream segment. A full enumeration and evaluation of all possible mappings is not tractable. Instead, the GM must utilize a guided search to find a solution.

We use a greedy search inspired by A^* [11] to quickly search for the best solutions. A^* search is a breadth-first, guided search algorithm that is guaranteed to find an optimal solution. An A^* search requires three components, an incremental search space, a non-decreasing cost function $f()$ and a heuristic function $h()$. The incremental search space allows the search to take place in a sequence of steps. For instance, a graph search is an incremental search space. For each edge traversed, the search has proceeded by one more step. A non-decreasing cost function $f()$ assures that for each step in the search, the cost will increase. In the graph setting, each edge would have a non-negative cost to insure that the cost will never decrease. Finally, the heuristic function $h()$ estimates the cost from the present position in the search space to a goal state. $h()$ is restricted in that it can never be higher than the true cost from the present location to the goal. In a graph setting, the heuristic function is commonly the Euclidean distance to the goal node. To conduct a search, A^* maintains a set of partial searches \mathcal{A} . New partial searches are added whenever an incremental search has to make a decision. For each step of the A^* search, the algorithm incrementally advances the partial search $a \in \mathcal{A}$ where $\forall_{b \in \mathcal{A}} f(a) + h(a) \leq f(b) + h(b)$. In other words, the A^* algorithm chooses the partial search that is most likely to lead to a solution. The search is guaranteed to find the optimal answer. However, a search improved by A^* is still exponential unless the error of the heuristic grows no faster than the logarithm of the true cost [17].

Our A^* inspired algorithm requires the same three components, an incremental search space, a non-decreasing cost function $f()$ and a heuristic function $h()$. Unlike A^* , our algorithm is depth-first search to conserve memory. The GM sequentially decides each frame-stream segment mapping starting at the first frame $v.b$ and ending at the last frame $v.e$. This forms the incremental search space. $f(p[])$ is the non-decreasing cost function where $p[]$ represents the array of all the stream segments that have been assigned to frames. This cost function can be calculated by summing the costs, $c()$, incurred in each time slot from time t_0 to $v.t + (v.e - v.b)$. The heuristic function $h(x)$ is the cost of the solution found by the simplified search algorithm for request $v' = (x + 1, v.e, v.t + (x + 1 - v.b))$. In other words, the heuristic function sums the lowest cost mapping for all non-assigned frames without considering that two frames could be generated in the same time slot. If two frames are generated at the same time, the total cost would be greater than if they were produced separately. Recall that the cost function $c()$ sums the square of each resource's utilization. Therefore, if we assume that no frames were generated at the same time, we can produce an estimate that will never exceed the true cost.

The pseudocode for the `FullSearch` algorithm is located in Figure 3. The GM picks the stream segment that can generate frame x with the lowest cost, ignoring any possible interactions with other frames (lines[5-7]). If the cost of the satisfied frames $f(p[])$ plus the heuristic cost $h(x)$ to satisfy the rest of the frames is less than the lowest cost c_l (line 10), the GM recursively solves for the next frame (line 13). Otherwise, the GM selects the next lowest cost stream segment (line 4). When there are no more stream segments available, the GM backtracks and repeats the process with the previous frame.

```

FullSearch( $p[], c_l, x$ )
1   $\mathcal{S}' := \mathcal{S}$ 
2  Remove all  $a \in \mathcal{S}'$  from  $\mathcal{S}'$ 
3    if  $(x < a.b)$  or  $(a.e > x)$ 
4  while ( $\mathcal{S}'$  not Empty)
5    choose  $s \in \mathcal{S}'$  such that  $\forall_{a \in \mathcal{S}'}$ 
6       $c(R(s.t - (x - s.b) + s.\vec{r}) \leq$ 
7       $c(R(a.t - (x - a.b) + a.\vec{r}))$ 
8    Remove  $s$  from  $\mathcal{S}'$ 
9     $p[x] := s$ 
10   if(  $f(p[]) + h(x) < c_l$ )
11     if  $(x = v.e)$ 
12        $c_l := f(p[])$ 
13     else FullSearch( $p[], c_l, x + 1$ )

```

Figure 3: Full Search Algorithm

5.2 Approximating Frame Space (Frame-Buckets)

The `FullSearch` algorithm does not take advantage of the fact that the best stream segment to generate frame x is very likely the best frame to service frame $x+1$. To exploit this trend, the frames of the requested video are grouped into *frame-buckets*. Instead of choosing a stream segment for each frame as in the exact algorithm, we now select a stream segment for each frame-bucket. Since the stream segment must satisfy all frames within a bucket, the algorithm only considers stream segments that can satisfy each frame in the bucket (see Figure 4(b)). By reducing the number of choices, the GM can more quickly come to a decision although not the greedy optimal one.

In this paper we examine two different bucketing algorithms (see Figure 4(a)). Equal spaced (Equi-spaced) buckets have an equal number n of frames in each bucket. Geometric spaced (Geo-spaced) buckets have m times more frames in each successive bucket, where m is a constant. For example an Equi-spaced bucketing algorithm with $n = 4$ would place 4 frames in each frame-bucket, while a Geo-spaced bucketing algorithm with $m = 2$ would fill the first 4 buckets with 1, 2, 4, 8 frames respectively. We will examine the effectiveness of both techniques in Section 7.2.

5.3 Approximating Resource Utilization (Time-Buckets)

In the exact algorithm, calculation of the values $f()$ and $h()$ depends on $\vec{R}(t)$. This means that the GM must calculate the cost for each value of t . To avoid recalculating $f()$ and $h()$ for each




Figure 4: Different bucket types: Buckets have 2 descriptors. One describes spacing of the buckets (a). Equi-spaced buckets each have the same number of frames in each bucket. Geo-spaced buckets geometrically increase (double) in size for each successive bucket. The second descriptor describes the resource type that the bucket is used to approximate over. Frame-buckets (b) are used to decrease the number of stream segments that are considered for the solution. If a stream segment does not provide all of the frames within a bucket, the stream segment is not considered in forming a solutions for those frames. Time-buckets (c) are used to approximate the resource utilization for large periods of time (measured in frames). The maximum utilization (for each resource type) serves as the approximation of the utilization across the entire bucket.

value of t , we use an enveloping technique [8] to represent resource utilization for larger periods of time. Time-slots are grouped into *time-buckets* in much the same manner as frames were grouped into frame-buckets in Section 5.2. For each time-bucket covering time-slots t_1 to t_2 , define $\vec{R}'(t) = (\max R_1(t) \dots \max R_k(t)); \forall t; t_1 \leq t \leq t_2$ as the maximum resource utilization during all of the time-slots within the bounds of the time-bucket (see Figure 4(c)). Since $\vec{R}'(t)$ has the same value for multiple values of t , values of $h()$ and $f()$ do not need to be recalculated for each time-slot but rather for each bucket.

6 Evaluation Environment

The GM is a scheduler for AMPS, a modular multimedia platform with stream reuse. To be effective, the GM must quickly schedule clients and provide the schedule to the system. The GM only needs to be fast enough to maintain the maximum throughput of the system without interfering with the system's performance. We use a simulation environment to determine if the GM can meet these requirements and discover any resource savings that the GM can offer. The GM is configured to operate as the scheduler for a forwarding proxy, the streams are not actually sent, but resource utilization is recorded. The set of stream segments \mathcal{S} is generated as detailed in previous work [4]. The GM accepts requests from the simulator and returns a schedule for each request. The GM must schedule all portions of the stream *by* the playback time, but the GM may schedule portions of the stream *before* their playback time.

The proxy has three stream modules: a network reception module that receives streams, a network transmission module that sends streams out, and a time shifting module that can buffer streams in memory until their playback time. The proxy also has 3 resources: memory, which is used by

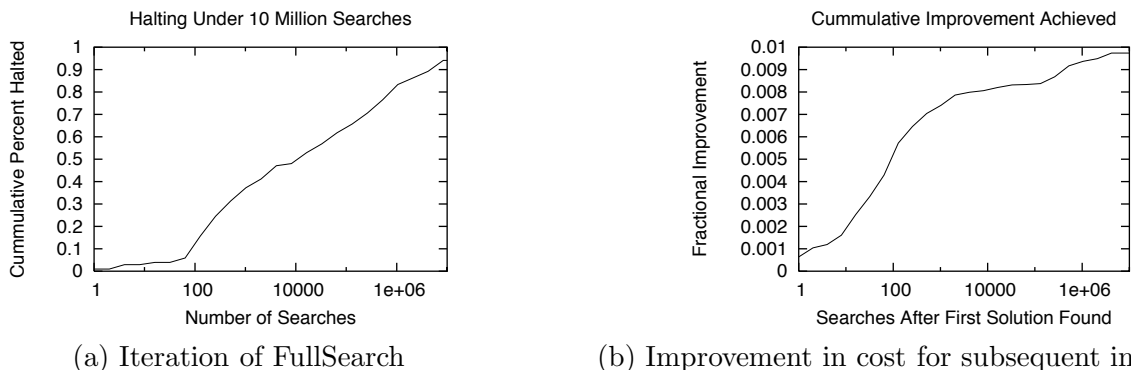


Figure 5: Baseline Search Characteristics

all three modules, server bandwidth, which is used by the network reception module (between the server and the proxy), and client bandwidth, which is used by the network transmission module (between the proxy and the client). The proxy has the average amount of each resource needed by a simple forwarding proxy to satisfy all requests. A simple forwarding proxy fetches each stream and forwards the stream to each requesting client. Note that with this amount of resources, it is unlikely that the simple forwarding proxy could service every arrival, as there is some irregularity in the arrival process.

We make use of two simulator settings: *Clip* and *Movie*. In the *Clip Setting*, the simulator generates requests for a single two minute video clip. In the *Movie Setting*, the simulator generates requests for a single one hour video. In both settings, the simulator generates requests using a Poisson arrival process with a mean of one request per second. The simulation is conducted for 4 times the length of the video type requested in order to reach steady state. Steady state occurs soon after the simulation has run for the length of the video. It is at this point that the system has (on average) a constant number of clients. All experiments were conducted on a Condor cluster with machines ranging from 2.4 GHz to 3 GHz processors.

7 Experiments

7.1 Baseline Search Characteristics

We conducted an experiment using the *Clip* setting to test the number of times line 10 of the `FullSearch` algorithm (Figure 3) is called during a search. For each request, the GM must search for the best solution. During the search the GM evaluated solutions until the `FullSearch` algorithm (see Section 3.3) identified it had found the lowest cost solution, or the algorithm had evaluated the cost of 10 million solutions. The CDF of when each search halted is presented in Figure 5(a). Five percent of the searches do not terminate before 10 million solutions have been evaluated. Each

Frame-Buckets	Time-Buckets								
	1s	5s	10s	15s	20s	30s	40s	60s	120s
1s	.7599	.7768	.7474	.6996	.6759	.6409	.6183	.5727	.7743
5s	.9060	.8622	.8067	.7804	.7614	.7215	.6805	.6448	.7219
10s	.9087	.9093	.9095	.8450	.7925	.7558	.7299	.6934	.7021
15s	.9096	.9082	.9077	.9083	.8746	.7987	.7447	.7051	.7189
20s	.9062	.9065	.9057	.9068	.9068	.8397	.8010	.7161	.6919
30s	.9052	.9052	.9052	.9054	.9065	.9052	.8716	.8099	.7192
40s	.9050	.9052	.9054	.9059	.9050	.9060	.9065	.8776	.7416
60s	.9050	.9052	.9050	.9050	.9050	.9050	.9052	.9050	.8328
120s	.9042	.9042	.9042	.9042	.9042	.9042	.9042	.9042	.9042

Table 1: Throughput achieved using Equi-spaced time- and frame-buckets

of these non-halting searches executed over 10 hours before the GM terminated the search. The returns of searching for longer periods of time are decreasing. The average improvement in cost after the first solution is shown in Figure 5(b). On average, the final solution’s cost is no more than 1% better than the first solution found. Searching for the solution with the lowest cost takes too long and provides little improvement. Our future experiments will only evaluate 1000 solutions to avoid long search times.

7.2 Evaluating Approximations

After examining the baseline behavior, the need for approximation is apparent. In this section we conduct a series of experiments to determine how well the approximations listed in Sections 5.2 and 5.3 perform. The results of each experiment are average of 10 runs. Time- and frame-buckets are either Equi-spaced or Geo-spaced buckets. When using Equi-spaced buckets, we will report the size of the bucket in seconds. Frame-buckets contain frames but only one frame is generated each second. When using Geo-spaced buckets, the first bucket is 1 second (frame) long. Each additional bucket is m times the size of the last bucket. So for a value of $m = 2$, the next three buckets would be 2, 4 and 8 seconds long respectively. We carefully chose values of m to insure that the sum of the frames in each bucket equals the number of frames in the Clip video.

Frame-Buckets	Time-Buckets								
	1s	5s	10s	15s	20s	30s	40s	60s	120s
1s	8130	2540	221	99.0	98.9	56.8	48.4	21.9	224
5s	196	71.0	38.1	34.8	25.5	19.7	20.0	14.1	36.7
10s	31.4	16.4	15.0	19.0	14.6	14.2	10.5	8.58	7.22
15s	31.5	12.6	13.1	10.2	13.3	9.79	9.87	6.45	4.68
20s	45.7	11.7	9.59	9.72	8.93	10.9	7.47	6.05	3.80
30s	32.2	10.6	8.71	9.58	8.32	8.81	8.15	5.09	3.55
40s	37.6	10.4	8.62	9.52	7.81	8.83	7.51	5.11	2.54
60s	36.4	10.6	9.01	8.50	9.66	9.45	9.01	9.18	3.67
120s	23.4	7.41	6.18	6.81	5.51	5.44	5.34	6.25	6.10

Table 2: Execution time in ms using Equi-spaced time- and frame-buckets

7.2.1 Equi-Spaced Buckets Using The Clip Setting

The first experiment examines the behavior of using Equi-spaced time- and frame-buckets under the Clip setting. The size of the buckets range from 1 second to 120 seconds, the size of the requested stream. The achieved throughputs are recorded in Table 1 and the execution times are recorded in Table 2.

This experiment exhibits several interesting results. First, we note that when the frame-bucket is 120s long the proxy behaves like a simple forwarding proxy that can only request streams from the server and relay them to the client. Because of natural fluctuations in the arrival process, 100% throughput is not possible and only 90.42% of requests are satisfied.

Second, when the time- and frame-buckets are 1s long, the approximation is the closest to the exact solution. The throughput is significantly less than that of a simple forwarding proxy. We examined the solutions generated by the search algorithm and discovered two trends. One, the solutions were composed of many small stream segments each 1-2 frames long. Two, many of these stream segments (sometimes as many as seven) were scheduled in the same time slot. Recall that the cost $c()$ is quadratic. The cost to generate two frames in the same time slot is greater than the cost to generate the frames in separate time slots. The search algorithm schedules frames in the same time slot only when the cost savings is greater than the cost penalty. In practice, the search algorithm finds a smaller increase in cost to schedule short stream segments compared to the savings gained from branching the stream. By scheduling frames at the same time, the proxy experiences short transient spikes in the resource utilization. These spikes in resource utilization only span 1-2 seconds. When the spike reaches 100% utilization, the proxy cannot schedule any streams during that time period. This shows that the greedy solution is not optimal for the global scheduling problem.

Frame-Buckets	Time-Buckets				
	1.533	2.39	3	4.6	10.6
1.533	.9064	.8119	.8000	.7581	.7363
2.39	.9075	.9050	.8728	.8081	.7742
3	.9068	.9085	.9050	.8795	.8165
4.6	.9067	.9072	.9075	.9067	.8617
10.6	.9060	.9055	.9075	.9091	.8998

Table 3: Throughput achieved using Geo-spaced time- and frame-buckets

Third, when the frame-buckets are smaller than the time-buckets, the throughput is poor. When the frame-buckets are as long as or longer than the time-buckets, we see throughput that matches or exceeds the throughput of the simple forwarding proxy (for lengths greater than 5 seconds). The greater throughput is due to the algorithm taking advantage of small dips in the resource utilization to schedule streams ahead of time. The proxy satisfies that client more quickly and is able to satisfy the next client as those resources are freed. However, this event occurs rarely and does not significantly improve the throughput.

Last, the execution time to find a solution decreases as the size of the buckets increase. Larger time-buckets decrease the execution time to calculate values of $h()$ and $f()$. Larger frame-buckets reduce the search space and thus make the search easier. Frame-buckets also produce a feedback effect. Each stream segment scheduled by the algorithm can be branched and reused to satisfy a later request. Fewer stream segments are scheduled when larger frame-buckets are used, and fewer choices are available in the future. This breaks down when the time-buckets are 120s long, in particular with 1s long frame-buckets. Since the time-bucket (only one) is the size of the stream, the system is using the peak resources utilization for the calculations of $h()$. As a consequence, the heuristic cost to produce a frame for any branched stream segment is the same. Since the costs are identical the guided search is not able to discard any of them, and the GM must evaluate each of them. In general, the approximations decrease the running time of the search.

7.2.2 Geo-Spaced Buckets Using The Clip Setting

This experiment examines the behavior using Geo-spaced time- and frame-buckets using the Clip setting. The size of each successive bucket is m times greater than the last. The achieved throughput is recorded in Table 3 and the execution time is recorded in Table 4.

The results of the Geo-spaced experiments are similar to the Equi-spaced results. When the frame-buckets are smaller than the time-buckets we see poor throughput. When the frame-buckets are as long as or longer than the time-buckets, we see throughput that matches or exceeds the throughput of the simple forwarding proxy. The execution times when using the Geo-spaced buckets decrease

Frame-Buckets	Time-Buckets				
	1.533	2.39	3	4.6	10.6
1.533	13.5	6.12	4.58	3.34	2.70
2.39	8.55	7.65	4.05	3.01	1.94
3	7.54	6.66	5.39	2.95	1.91
4.6	6.70	5.62	5.17	4.26	1.84
10.6	5.99	5.12	4.51	3.84	2.29

Table 4: Execution time in ms using Geo-spaced time- and frame-buckets

Frame-Buckets	Time-Buckets				
	1.533	2.39	3	4.6	10.6
1s	.6848	.6228	.5983	.6176	.6422
5s	.7574	.6527	.6452	.6313	.6405
10s	.8038	.6848	.6601	.6473	.6643
15s	.8515	.7199	.6847	.6499	.6671
20s	.8316	.7482	.7135	.6728	.6745
30s	.8878	.8351	.7335	.7258	.6889
40s	.9059	.8420	.8613	.7689	.7337
60s	.9054	.8970	.8792	.8478	.7574
120s	.9042	.9042	.9042	.9042	.9042

Table 5: Throughput achieved using Geo-spaced time-buckets and Equi-spaced frame-buckets

as the size of the buckets increases. However, there is a (small) speedup when using Geo-spaced buckets over Equi-spaced buckets in the Clip setting.

7.2.3 Combination Spaced Buckets Using The Clip Setting

We performed additional experiments to test how using one bucketing technique for time-buckets and different bucketing technique for frame-buckets affect throughput using the clip setting. When using Geo-spaced time-buckets and Equi-spaced frame-buckets (see Table 5), the system fails to achieved good throughput (over 90%) for all but large frame-buckets. When using Equi-spaced time-buckets and Geo-spaced frame-buckets (see Table 6), the trends in throughput are similar to the trends seen when using only one bucketing technique for both types of buckets. However, the trend is not as clearly defined. Instead we will focus further attention on only using one bucketing technique for each experiment.

Frame-Buckets	Time-Buckets								
	1s	5s	10s	15s	20s	30s	40s	60s	120s
1.533	.9087	.9014	.8965	.8897	.8861	.8722	.8454	.7959	.7770
2.39	.9067	.9042	.9000	.9001	.8991	.8914	.8903	.8598	.8048
3	.9069	.9051	.9039	.8986	.8994	.8984	.8894	.8759	.8245
4.6	.9062	.9054	.9037	.9037	.9040	.9004	.8999	.9001	.8394
10.6	.9055	.9054	.9052	.9039	.9041	.9040	.9043	.9047	.8784

Table 6: Throughput achieved using Equi-spaced time-buckets and Geo-spaced frame-buckets

Bucket Size	5 min	7.5 min	10 min	15 min	20 min	30 min	60min
Segments	12	8	6	4	3	2	1
Throughput	.9832	.9829	.9828	.9827	.9830	.9827	.9828
Exec. Time	823	697	445	362	308	313	445

Table 7: Hour long experimental results using identically sized, Equi-spaced time- and frame-buckets

7.2.4 Hour Length Experiments

Our next pair of experiments measure the throughput and execution times of longer streams. We conduct each experiment using the Movie setting. In the first experiment, we examine the performance of using Equi-spaced time- and frame-buckets. In the second experiment we examine the performance of using Geo-spaced time- and frame-buckets. For both experiments the time- and frame-bucket are the same length, and the parameters are only reported once. The results of using Equi-spaced buckets are presented in Table 7. The throughput achieved by a simple forwarding proxy is 98.28%. We see this throughput when the time- and frame-buckets are both 60 minutes long. The throughput for other bucket sizes are similar, but the execution time decreases for larger buckets (except for the 60 min buckets). The results from using Geo-spaced buckets are presented in Table 8. The throughput is similar to the throughput achieved using Equi-spaced buckets. However, the execution times experienced when using Geo-spaced buckets are up to an order of magnitude lower than using Equi-spaced buckets.

Multiplier	1.447	1.975	2.34	3.045	4.925	7.48	14.98	59.5
Segments	20	12	10	8	6	5	4	3
Throughput	.9839	.9836	.9834	.9836	.9832	.9829	.9829	.9818
Exec. Time	433.7	265	213	170	127	91.3	37.3	16.0

Table 8: Hour long experimental results using identically sized, Geo-spaced time- and frame-buckets




Figure 6: Average Peak Resource Utilization using identically sized, Equi-spaced time- and frame-buckets

7.3 Effectively Using Resources

In this section, we report the average resource consumption using the Clip setting for each of the three resources: client bandwidth, server bandwidth and memory. For each bucketing technique used, both time- and frame-buckets have the same size. For each experiment, we provide the average resource consumption per time-bucket. The size of each bucket is determined by the bucketing technique and the parameter used. The resources utilization using Equi-spaced buckets are found in Figure 6, and the resource utilization for Geo-spaced buckets are found in Figure 6. In reading these graphs the most important data is the utilization of the first bucket. In all graphs the resource utilization drops off after the first bucket and the utilization of the first bucket limits the ability to service new requests.

The client bandwidth consumed is similar for both bucketing techniques and for all parameter values (Figure 6(a) and Figure 7(a)). This resource is tied directly to the proxy’s throughput. As each of the settings achieve throughput around 90%, we do not expect to see any difference in client bandwidth consumption.

We see greater differences in the manner in which server bandwidth and memory are consumed (Figures 6(b,c) and Figures 7(b,c)). When using Equi-spaced buckets, the resource utilization decreases when *smaller* buckets are used. When using Geo-spaced buckets, we observe the opposite trend. The resource consumption decreases when *larger* buckets are used. Geo-spaced buckets provide a double win since using larger buckets decreases both the resources consumed and the execution times needed to satisfy a request.

8 Discussion

Given any system that performs one action (such as reading a file from disk) with a static workload, we are confident that a skilled team of programmers can develop the software necessary to meet that challenge. However, the GM is useful for any multimedia system that exhibits one or more

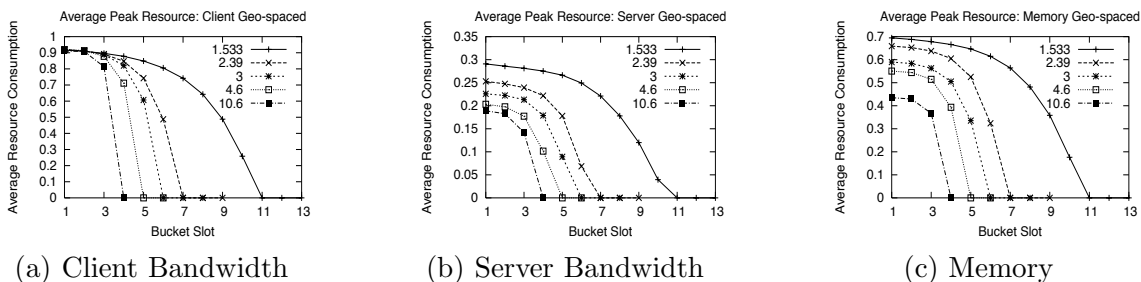


Figure 7: Average Peak Resource Utilization using identically sized, Geo-spaced time- and frame-buckets

of the following traits: 1) contains a large set of stream modules in which multiple sequences can satisfy the same request; 2) concurrently provides a stream to multiple clients using different services and different starting times; 3) offers a large and changing set of services; and/or 4) is resource-constrained. When there are multiple sequences of stream modules that can satisfy a request, the GM can determine which is best, given the present resource usage. When a server or proxy must serve the same stream, but offer different services and different starting times to different clients, the GM determines those portions of existing streams that can be reused, by buffering within the system or by providing more than one stream to the client. When a system needs to offer a large and changing set of services, the GM can help the system avoid recompilation. The GM determines the best way to use new modules in coordination with old modules to satisfy requests. This ability avoids costly rewrites when new modules become available. When resources are constrained, the GM can determine how to utilize resources to service more requests.

For any benefit, we can expect to pay a cost. The time required to schedule streams using the GM is likely to be higher than the time needed for a straightforward system. In this paper we provided measurements to determine the speed with which the GM can schedule and have found the cost per client to be low, on the order of 10s of milliseconds for hour long streams with an average of 3600 simultaneous clients. Recall that each stream used to serve the existing 3600 clients must be evaluated to determine if they can be reused for the new client. Workloads with less sharing require less processing. But, is 3600 clients at reasonable limit? The AMPS proxy [20] supported a maximum of 710 clients using 300k streams using a 2 GHz processor. Performance studies of Microsoft and RealNetworks servers [12] exhibited maximum throughput rates for 500k streams well under 1000 simultaneous clients using dual 700 MHz processors. Since these experiments were conducted on different servers and under different conditions they cannot be compared directly, but used to illustrate arrival rates that entry level multimedia servers can support. The overhead of using the GM on these systems is minimal, but offers the dual benefits of providing more services and more efficient use of resources.

References

- [1] Soam Acharya and Brian Smith. Middleman: A video caching proxy server. In *Proc of NOSSDAV*, 2000.
- [2] J. M. Almeida, D. L. Eager, and M. K. Vernon. A hybrid caching strategy for streaming media files. In *Proc. of MMCN*, January 2001.
- [3] Elan Amir, Steven McCanne, and Hui Zhang. An application level video gateway. In *Proc. ACM Multimedia*, November 1995.
- [4] Michael K. Bradshaw, Jim Kurose, Lela Jane Page, Prashant Shenoy, and Don Towsley. A reconfigurable, on-the-fly, resource-aware, streaming pipeline scheduler. In *Proc. of MMCN*, January 2005.
- [5] David W. Brubeck and Lawrence A. Rowe. Hierarchical storage management in a distributed vod system. *IEEE MultiMedia*, 3(3):37–47, 1996.
- [6] Michael W. Carter and Craig A. Tovey. When is the classroom assignment problem hard? *Oper. Res.*, 40(S1):28–39, 1992.
- [7] A. Dan and D. Sitaram. A generalized interval caching policy for mixed interactive and long video environments. In *Proc. of MMCN*, 1996.
- [8] Victor Firoiu, Don Towsley, and Jim Kurose. Efficient admission control for edf schedulers. In *Proc. IEEE INFOCOM*, April 1997.
- [9] Steven D. Gribble and et al. The ninja architecture for robust internet-scale systems and services. *Journal of Computer Networks*, 35(4), March 2001.
- [10] GStreamer Team. GStreamer: open source multimedia framework. <http://www.gstreamer.net>.
- [11] P. E. Hart, N. J. Nilsson, and B. Raphael. A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics*, pages 100–107, 1968.
- [12] KeyLabs. Helix universal server from realnetworks comparative load test, May 2002. www.keylabs.com.
- [13] Zhuoqing Morley Mao, Hoi sheung Wilson So, and Byunghoon Kang. Network support for mobile multimedia using a self-adaptive distributed proxy. In *Proc of NOSSDAV*, 2001.
- [14] Microsoft Inc. Directshow. <http://msdn.microsoft.com>.
- [15] Gerard J. M. Smit et al. Reconfiguration in mobile multimedia systems. In *Proc. of PROGRESS workshop*, October 2000.

- [16] Z. Miao and A. Ortega. Proxy caching for efficient video services over the internet. In *Packet Video Workshop*, April 1999.
- [17] J. Pearl. *Heuristics: Intelligent Search Strategies for Computer Problem Solving*. Addison Wesley, Reading Massachusetts, 1984.
- [18] Edward J. Posnak, Harrick M. Vin, and Robert G. Lavender. Presentation processing support for adaptive multimedia applications. In *Proc. of MMCN*, January 1996.
- [19] R. Rejaie, M. Handley, H. Yu, and D. Estrin. Proxy caching mechanism for multimedia playback streams in the internet. In *Proc. of the 4th International Web Caching Workshop*, March 1999.
- [20] X. Zhang and et al. AMPS: A flexible, scalable proxy testbed for implementing streaming services. In *Proc of NOSSDAV*, Kinsale Ireland, June 2004.