EVALUATING THE EFFECTIVENESS OF AUTOMATIC PVR MANAGEMENT

Ketan Mayer-Patel

University of North Carolina, Chapel Hill
Dept. of Computer Science
kmp@cs.unc.edu

Wesley Miaw

University of North Carolina, Chapel Hill
Dept. of Computer Science
wesley@cs.unc.edu

ABSTRACT

A model for evaluating the effectiveness of automatic recording of television programs by digital personal video recorders (PVRs) is presented. The model is used to evaluate the tradeoff between manual management of recording programs to a PVR archive and automatic management of recording. We show that a tradeoff exists between manual management in which the utility of a program is perfectly known but user awareness of available programs is limited versus automatic management in which utility estimates contain error but awareness is perfect. Experiments with the model show that the shape of this tradeoff is most governed by the shape of the distribution of user utility. As the percentage of programs with high user utility decreases relative to the average, the more effective automatic recording is likely to be despite errors in utility estimation. The shape of this tradeoff, however, is highly inelastic. Thus, improving utility estimates will not make automatic recording more effective if user awareness is sufficiently high.

1. INTRODUCTION

We are now poised on the brink of a new day for television. High-definition television is rapidly gaining acceptance and the Internet has become an alternate means for distributing video content. Advances in compression and streaming technologies have made video-on-demand possible. Furthermore, digital cable television systems increasingly represent the norm for television distribution and the number of channels available to the consumer has vastly increased.

As we move toward the paradigm of digital television, it is not surprising that one of the first appliances to be developed using much of this new technology is a digital version of the ubiquitous and indispensable VCR. A digital personal video recorder (PVR) is an information appliance that, like the VCR, provides for time-shifted viewing and archiving of television programs. The PVR, however, extends this capability in a number of important ways. First, the PVR is capable of playing and recording at the same time. This allows the user to access archived programs while recording new ones. Second, the PVR provides instant and easy access to an entire archive of television programs. Current PVRs are capable of storing anywhere between 30 and 300 hours of programming. Third, PVRs are generally integrated with a meta-data service, allowing users to search for and schedule recordings based on descriptions of content.

These extended features allows PVRs to do what VCR's never could: proactively and speculatively record programs on behalf of the user without explicit programming to do so. Several PVRs currently on the market (e.g., TiVo [1], ReplayTV [2], UltimateTV [3]) already provide this functionality in some form or another. Used in this way, a PVR operates less like a VCR and more like an automatically managed local cache wherein video streams are stored based on the likelihood that the user will want to access these streams at some time in the future.

One obvious shortcoming of automatic PVR management is that the algorithm used to determine user interest in a program may by inaccurate. Space in the archive is then wasted on programs that the user has no use for. On the other hand, automatic PVR management has the advantage of being able to scan every program broadcast on every channel at any time of day. We would not expect a person to be as aware of everything that is being broadcast. In this paper, we investigate this tradeoff between a person's ability to identify programs given imperfect knowledge about what is available and perfect knowledge about the utility of a program versus an automatic recording algorithm with perfect knowledge about what is available and imperfect knowledge about utility.

The main contributions of this paper are:

- An abstract PVR model which allows us to simulate different automatic recording algorithms and human recording behaviors.
- Experiments which explore the tradeoff between awareness and utility estimation.

Our initial experiments show that the shape of this tradeoff is most governed by the shape of the distribution of user utility for a particular program. Furthermore, as the percentage of programs with high user utility decreases relative to the average, the more effective automatic recording is likely to be despite errors in utility estimation. In other words, as the number of programs overwhelms the viewer, finding high utility programs becomes harder to do given limited awareness and thus automatic recording can prove effective. The shape of this tradeoff, however, is highly inelastic. Thus, increasing user awareness by only a few percentage points will generally make up the difference in the expected performance between automatic recording with 50% utility estimation error and 5% utility estimation error.

The rest of this paper is organized in 5 sections. Section 2 motivates our approach. Section 3 describes our abstract PVR model. Results from initial experiments with the model are given in Section 4. Section 5 discusses related work. The paper is summarized and our plans for future work are described in Section 6.

2. MOTIVATION

Our work is motivated by two anticipated characteristics of digital television in the future:

- **The amount of available content is overwhelming.** Today's analog cable distribution networks contain anywhere between 50 and 150 channels. As digital systems (both cable and satellite) become more widespread, this number can be expected to grow to several hundred. This amounts to potentially between 500 and 1000 hours of video programming made available to the user daily. Finding programs of interest in this sea of information will become increasingly difficult.

- **Broadcast distribution is effective and necessary.** The size of video program data (roughly 1GB/hour) makes broadcast distribution practical and necessary. Making use of off-peak viewing hours to push the most popular content closer to the user will be important for reducing server and network load in video on demand systems.

In this work, our goal is to explore the tradeoff between two policy extremes which we believe represent endpoints of a spectrum. On one extreme is a policy that mimics user behavior with no automatic recording whatsoever. We call this the *manual* policy. On the other extreme is a policy that automatically manages all recording. We call this the *automatic* policy. The tradeoff in question relates to the issue of user awareness and error in judging the utility of a particular video program.

The manual policy is effective because the user is perfectly aware of the utility of any given program when it is initially broadcast. However, the manual policy is limited in its awareness of what is being broadcast. Thus, the manual policy is unable to detect some high-utility programs which then go unrecorded, recording instead programs that may be of lesser utility. The automatic policy has the advantage of being aware of every program broadcast and thus is able to consider all programs for recording. The automatic policy is burdened, however, by imperfect knowledge about the utility of any particular program to the user. This imperfect knowledge stems from errors and inaccuracies in whatever algorithm is used to determine utility. Here we consider this as an abstract function and only characterize the error without attaching the work to any specific algorithm.

Our hypothesis is that there is a tradeoff between awareness and utility error such that the manual policy will outperform the automatic policy given sufficiently high awareness and/or utility error. Similarly, we expect the automatic policy to dominate given sufficiently low awareness and/or utility error. The purpose of the experiments presented here is to confirm this hypothesis and map the shape of this tradeoff as it relates to a number of different parameters such as the number of programs available for recording, user consumption of video data, and the decay of utility as a function of time.

3. PVR MODEL

We conducted our experiment by simulating viewer and PVR behavior using a model designed to closely resemble the capabilities of common consumer PVRs and the television industry's current distribution system. We did not include some of the more advanced PVR capabilities that have been recently added such as multiple tuners. Our model has five independent variables.

1. **Channels**
   The number of channels, which dictates the number of programs available each half hour. We experimented with 100, 150, 200, 250, and 300 channels.

2. **Consumption**
   The number of half hour blocks the viewer watches each day. Afterwards, the watched programs are removed from the PVR. We experimented with consumption rates of 2, 4, 8, 12, and 16 half hour blocks consumed per day.

3. **Content Utility Distribution (CUD)**
   Each program's utility to the viewer falls within the range $[0, 1]$ and is randomly chosen from a gaussian curve described by $e^{-z^2(x-0.5)^2/2}$ where x is a random number between 0 and 1 and z defines the sharpness of the distribution. We experimented with z values of 8, 10, 20, 30, and 40. A sharper distribution will result in fewer shows of high utility than a less sharp distribution.

4. **Decay rate**
   Every day a program remains unwatched, we reduce its utility by a multiplicative factor. We experimented with decay rates of 1.000, 0.975, 0.950, 0.925, and 0.900.

5. **Policy**
   The policy is used to decide which programs to store and replace in the PVR. The manual policy recognizes a program's actual utility but is only aware of a small percentage of all programs. The automatic policy estimates a program's utility within some margin of error but is aware of all programs. The automatic policy error distribution is uniform. We experimented with manual awareness percentages ranging from 1% to 15% and automatic error margins of ±0.5 to ±0.50 in increments of 0.05.

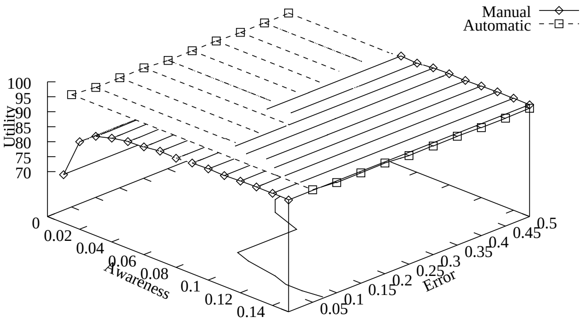

Figure 1: Manual versus automatic mean utility.

Each experiment simulates 180 days with a PVR storage capacity of 100 half hour programs. Each half hour, the program with highest utility among those the policy is aware of is added to the PVR. Programs of lesser utility are removed from the PVR to make room if necessary. After each simulated day, the total utility of the stored programs is computed and the highest utility programs "consumed" by the user are removed. The first simulated week is ignored to allow the PVR to reach capacity.

4. RESULTS

The manual policy's mean daily utility increases with awareness while the automatic policy's mean daily utility decreases with its margin of error. This tradeoff is very inelastic; a slight increase in awareness can make a large difference while the margin of error has very little effect on the performance of the automatic policy.

Figure 1 illustrates this tradeoff. The manual policy and automatic policy are compared for 200 channels, a consumption rate of 8 half hour programs per day, CUD z=10, and utility decay factor of 0.975. The intersection of the two surfaces is projected below onto the awareness-error plane. The automatic policy outperforms the manual policy at those points on the awareness-error plane between the origin and the intersection. The manual policy outperforms on the other side of the intersection.

The impact of different CUDs, decay rates, number of channels, and consumption rates will be presented by comparing the surface intersections of each simulation. In all of the following experiments, unless otherwise specified, the number of channels is 200, the CUD z factor is set to 10, consumption is set to 8 half hour programs per day, and the decay rate is set to 0.975.

4.1. Channels

Since awareness is defined as a percentage of the total number of available programs, the manual policy does better as the number of channels increases. This stands to reason as the same percentage of more channels gives the manual policy more choices each half hour. In Figure 2, the number of channels varies from 100 to 300.

4.2. Consumption Rate

Figure 3 shows that increasing the consumption rate marginally improved the performance of the automatic policy in comparison to the manual policy. In this experiment, consumption

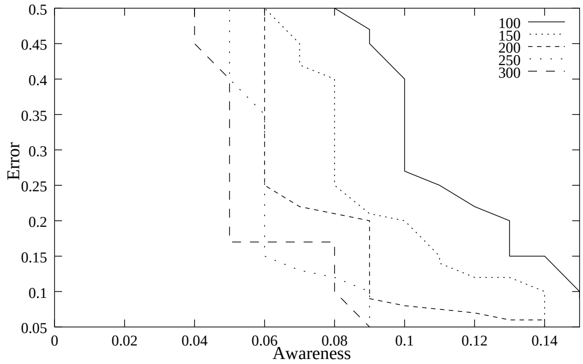

Figure 2: The effect of varying number of channels.

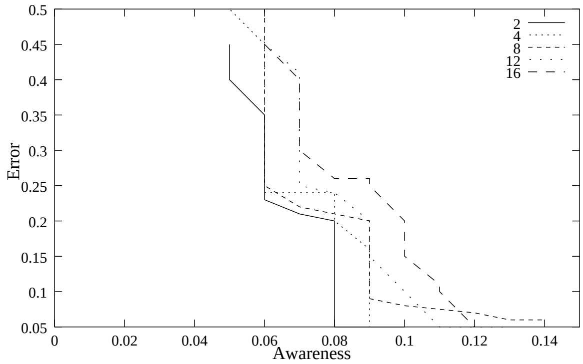

Figure 3: The effect of viewer consumption rate.

is varied between 2 and 16 half hour programs per day. A higher consumption rate means the policy must replace more programs of high utility each day. This is easier to accomplish for the automatic policy since it has perfect awareness.

4.3. Content Utility Distribution

As expected, as the content utility distribution curve sharpens, the manual policy needs higher awareness to outperform the automatic policy, even when the automatic policy has a high margin of error. Figure 4 displays the intersection lines for $z \in \{8,10,20,30,40\}$.

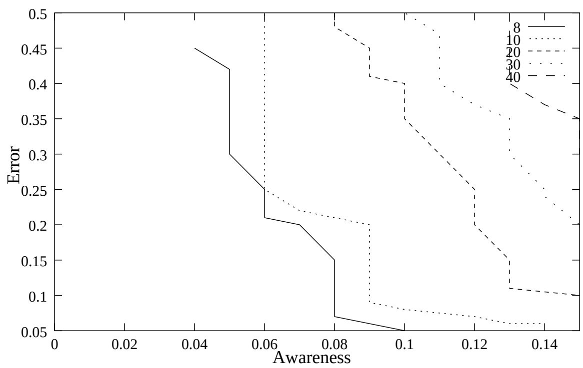

Figure 4: The effect of varying content utility distributions.

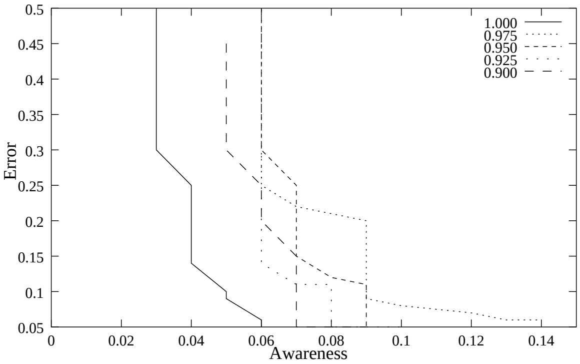

Figure 5: The effect of varying utility decay.

4.4. Decay Rate

With no decay (decay rate 1.0), the manual policy does very well because it can accumulate programs of high utility over an arbitrarily long period of time. Since the viewer cannot consume as many programs as he is aware of, this will eventually result in a PVR storing only programs with utility equal to or very close to the maximum utility.

However, with decay rates less than 1.0 both policies must try to maximize total utility in as short a time period as possible. This is clear from Figure 5 where there is a wide separation between the intersection when decay rate is 1.0 and the intersections produced by the other four decay rates. It is notable that there is not much difference between the intersections of the four decay rates less than 1.0.

5. RELATED WORK

Because the introduction of the PVR as an information appliance is a relatively recent phenomena, there is not a body of previous work that directly addresses the issues raised in this paper Areas of research that are more indirectly related, however, include proxy caching of multimedia objects and collaborative filtering and recommender systems.

The goal of proxy caching of multimedia systems is to manage in-network caches of video data which may be useful to one or more members of a group of users in order to minimize the user of network and/or server resources. We can consider the PVR as acting as a proxy cache for just one end user. In [4], prefix caching was shown to be an effective technique for dramatically reducing server load in video-on-demand systems and hiding start up latency. Although we consider caching of whole video objects without video-on-demand, we can adapt our model to consider prefix caching with video-on-demand simply by reducing the expected size of the media objects. This work was extended in [5] to consider the effect of prefix caching with periodic broadcast. Other media caching systems that investigate cache space allocation issues and the impact on networking resources include [6, 7, 8] and [9].

Research into collaborative filtering and recommender systems has produced a number of techniques for automatically estimating, with various degrees of user input and effort, the utility of a media object for a particular user. Our model is independent of a specific scheme and simply characterizes the error of this estimation. Collaborative filtering schemes correlate user interests with those of other users in order to identify media objects with high utility. Media type-independent work in this area includes [10, 11]. In [12], a television-specific recommender system is developed.

6. SUMMARY

In this paper, we investigated the tradeoff between a person's ability to identify high utility programs given imperfect knowledge of what is available and perfect knowledge about utility versus any automatic recording algorithm which would have perfect knowledge about availability but imperfect knowledge about utility. To do this, we developed a PVR model and simulated these two policies under a variety of conditions.

The most interesting result of our experiments is that the margin of error has very little effect on the performance of the automatic policy, while awareness has a large effect on the performance of the manual policy. Thus, if user awareness is sufficiently high, an automatic management policy will not be able to perform as well even if its utility estimation error is quite low.

Our experiments, however, do show that an automatic policy can be of great benefit to viewers who are only aware of a small percentage of the available programs or have very particular tastes, even when the margin of error is very high. Furthermore, the tradeoff between manual and automatic policies is heavily influenced by the content utility distribution while the consumption and decay rates do not have a significant impact.

7. REFERENCES

[1] http://www.tivo.com.
[2] http://www.replaytv.com.
[3] http://www.ultimatetv.com.
[4] S. Sen, J. Rexford, and D. Towsley, "Proxy prefix caching for multimedia streams," *Proc. of IEEE INFOCOM*, April 1999.
[5] Y. Guo, S. Sen, and D. Towsley, "Prefix caching assisted periodic broadcast: Framework and techniques to support streaming for popular videos," Tech. Rep. UM-CS-2001-022, University of Massachusetts, May 2001.
[6] S. Acharya and B. Smith, "Middleman: A video caching proxy server," *Proc. of NOSSDAV*, 2000.
[7] K.-L. Wu, P.S. Yu, and J.L. Wolf, "Segment-based proxy caching of multimedia streams," *Proc. of the 10th International WWWW Conference*, 2001.
[8] E. Balafoutis, A. Panagakis, N. Laoutaris, and I Stavrakakis, "Study of the impact of replacement granularity and associated strategies on video caching," *Cluster Computing special issue on networking*, To appear in 2002.
[9] C. Venkatramani, O. Verscheure, P. Frossard, and K.-W. Lee, "Optimal proxy management for multimedia streaming in content distribution networks," *Proc. of NOSSDAV*, 2002.
[10] B. Miller, J. Riedl, and J. Konstan, "Experiences with grouplens: Makeing usenet useful again," *USENIX Annual Technical Conference*, 1997.
[11] D. Goldberg, D. Nichols, and B.M. Oki, "Using collaborativ filtering to weave an information tapestry," *Communications of the ACM*, vol. 35, no. 12, December 1992.
[12] P. Baudisch, "Recommending tv programs: how far can we get at zero user effort?," *AAAI Workshop on Recommender Systems*, 1998.