

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

With high-definition television rapidly gaining acceptance and the Internet as an alternate means for distributing video content, we are entering a new era of television. 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 be 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

