Design and Evaluation of a Modern Switched Digital Video System

Adil Alsuhaim
School of Computing
Clemson University
Clemson, SC 29634
Email: aalsuha@clemson.edu

Zhaoyu Jing
School of Computing
Clemson University
Clemson, SC 29634
Email: zjing@clemson.edu

James Martin
School of Computing
Clemson University
Clemson, SC 29634
Email: jmarty@clemson.edu

***Abstract*—In this paper we implemented a switched digital video broadcast simulator to investigate the effect of using limited bandwidth on the perceived quality from the users' perspective. Our results show that it is possible to deliver a reasonable quality service even when allocating lower bandwidth to the content channels offered by a service provider. We think that further research into what ratio of bandwidth should be allocated to a given number of subscriber base.**

***Index Terms*—Switched Digital Video, Simulation, Spectrum Efficiency.**

I. Introduction

Switched Digital Broadcast (SDB) is a newer method of distributing video programming. Compared with traditional broadcast methods, it reduces spectrum requirements by taking advantage of the fact that not all program channels are being viewed by subscribers at the same time. Only a fraction of the channels are ever watched at given time, allowing video providers to offer more "virtual channels" than physically allowed by the spectrum. There are statistical gains achieved through over-provisioning of the available resources.

In a traditional video broadcast, the provider would stream all available channels towards users using frequency division multiplexing, and users can pick channels that they desire to watch with the remote control. This corresponds to picking up a frequency within the spectrum on which the channel is carried. Since not all the content channels are being watched, a part of the spectrum is wasted.

In a real-world SDB deployment, channels are only streamed when requested by one or more users, saving bandwidth and resource, and allowing more content channels to be offered, and improve the overall quality of the service. However, the actual spectrum savings depends on how the behavior of subscribers on how the watch TV, the popularity of delivered TV programs, streaming bit-rate composition and subscriber group size[1].

The practice of over-provisioning resources is not uncommon. Airlines oversell tickets in anticipation of no-shows. In which case, passengers might be blocked from boarding a plane. Airlines would ask travelers to forgo their seats in exchange for a generous compensation. We believe that in switched digital broadcast that there is more flexibility than other industries. A service provider would likely want to offer a quality service with the least resources, while minimizing blocking of those services.

II. Related Work

The development of the Internet services has opened the door for better opportunities to provide users with layers of services on top of a switched network. Proposals and anticipations of the future of switched digital videos had been published in the past two decades [12].

Research focuses on managing bandwidth in digital video broadcast to provide users with better quality content, improve overall satisfaction with the service while reducing bandwidth consumption. The research is influenced by how subscribers use digital video services. Request from a fixed number of users for web services follows is distributed according to zipf law [2]. In our paper, we assume that TV channel request, or channel popularity, follows Zipf law [9].

Managing video requests of different quality types can be modeled as a bin packing problem, where different bandwidth requirement are fitted within steams of fixed bandwidth capacities [1]. However, when it comes to TV service over IP Network, user's behavior would including switching between channels. `SimulWatch` [8] attempts to model such behavior, which achieved accurate modeling of the performance of IPTV systems. [13] proposes a control algorithm by formulating channel switching into an optimization problem, and then using dynamic programming to reach an optimal solution that minimize channel zapping time and overall bandwidth use. It uses Zipf law to rank channel popularity and Poisson distribution to model channel change request arrivals. A study on logs form a video-on-demand systems in [11] concluded that popularity of videos follows a Zipf law as well. However, it concluded that Poisson model overestimates the probability of large arrival groups. [14] use a predictive model to anticipate user's behavior to reduce the delay in channel switching. Other solutions to minimize channel switching time when the user is flipping up or down through TV channels by ordering them in an optimal manner [15].

The remaining sections of this paper are as follows: section 3 describes our system's overview. Section 4 discusses a set of experiments that we have conducted with our simulator

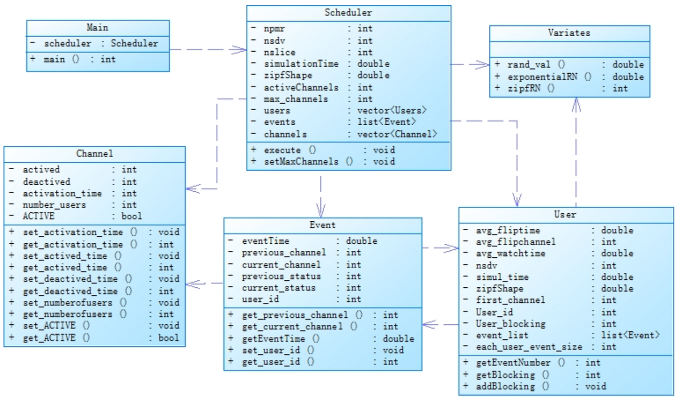

Fig. 1. Basic SDV system

. Section 5 is a discussion of the results. Section 6 is our conclusion, and recommendations.

III. Simulation System Overview

We implemented a high level, C++ simulator environment that mimics a switched digital broadcast deployment. It models subscribers (users), content channels and the bandwidth on which the channels would be transmitted to the end users. Users are classified according to their viewing behavior. A user is either flipping through channels or watching a certain channel for an extended period of time. Channels are ranked according to their popularity, and therefore the users of the system are more likely to view or flip through channels that are more popular than others. The bandwidth is represented by slices, an integer value that end users contend over. The higher the number of slices, the more bandwidth is available between users and the service provider. Higher resolution TV channels would require more slices than lower resolution channels. Figure 1 depicts the C++ class diagram for our simulator and how the components interact.

A. Users

We modeled users behavior into two categories. Users that exhibit a channel flipping behavior, and users that do not exhibit such behavior.

1) User with Flipping-Watching Behavior: In this category users are either watching a particular content channel or flipping through a series of channels, simulating how users browse through TV channel to look for an interesting program. In our simulation, a user belongs to either a `FLIPPING` or a `WATCHING` state. We implemented "events" to describe how users actions transition them between these two states as illustrated in Figure 2. No action is required when a user is in watching channel. Users either flip through channels, and then at the last flip, they would stop at a channel, effectively transitioning them to a watching state.

Based on test data obtained from two SDB trials in [9], the relationship between the total content channels and the number of content channels viewed concurrently can be modeled with Zipf distribution[16]. Zipf distribution refers to that, if we rank order all the content channels from the most popular to the least popular, the distribution implies that high-ranked content channels will be selected to watch extremely frequently, whereas low-ranked channels are rarely chosen. For our problem, we describe Zipf distribution for the probability of the ith channel being watched as follows:

$$P(i) = C\frac{i^{-\alpha}}{\sum_{i=1}^{n} i^{-\alpha}} \quad (1)$$

where C is a scaling factor, α is Zipf shape parameter, and n is the total number of content channels.
For this type of users, we start up all users into `FLIPPING` states, where each user flips through k channels before they enter a `WATCHING` state. After k flips, a user stops to watch a content channel for a relatively long period of time. After the watching time passes, the user would enter into a `FLIPPING` state again, flipping through channels and so on.

The number of channels each user flips through, the amount of time they spend before flipping away, and the average amount of time users remain at watching state are generated randomly from exponential distributions with a programmer-configurable mean values for each. In addition, the flipping events are classified into flipping increasingly, decreasingly or randomly. The increasing/decreasing behavior emulate a user pressing up or down on a remote control. The random flipping behavior emulates a user going through their list of favorite channels, and therefore, the next channel a user flips to is generated from a Zipf distribution.

2) User without Flipping Behavior: Compared with the first category, the other kind of user behavior does not exhibit any flipping behavior. This type of users emulates subscribers that use their Digital Video Record (DVR) to record programs at scheduled times. For this kind of behavior we define two states, `RECORDING` and `DORMANT`. A DVR user enters a `RECORDING` state at a scheduled time to record from a content channel. After the recording is done, the user enters the `DORMANT` state, in which it sleeps until the next scheduled event. The choice of channels is generated using Zipf distribution, and the watching and sleeping times are generated randomly from an exponential distribution with user-configurable means.

In a real world scenario, users may flip through channels and use the DVR feature of their set-top box. For simplicity, we assume that users that flip through channels and users that

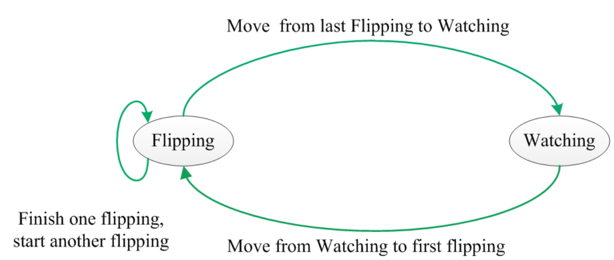

Fig. 2. Flipping User behavior

user DVR are two separate categories.

B. Channel slices

We use integer values to represent what we call ”slices”. Channel slices corresponds to the bandwidth division into multiple spectrum. This allows up to simplify the simulation into a spectrum efficiency problem. Multiple channel slices can be consumed one content channel. If the channel is of a higher quality video, it would require higher number of slices as detailed in Table I. When a user requests a channel, the number of available channel slices are reduced to represent bandwidth consumption by a subscriber based. If another user from the same subscribers' base requests a channel for which slices have been allocated, the user will get their request met without consuming any additional slices. However, if a user request a channel is request and there were no allocated slices for it, the simulator will consume more slices to satisfy that request. We may run into the situation where there are not enough slices available. In this case, we may choose to block the channel request, or choose to downgrade the quality of the requested channel. We made our simulator configurable to accommodate for both scenarios. The algorithm to downgrade the quality is detailed in algorithm 1, where the number of slices allocated for a channel is reduced to avoid blocking.

C. Content Channels

Content channels are provided for users to choose from. Users pick channels on basis of popularity provided by Zipf law. In our simulation, we have a data structure to hold the channels. The channels are available in all levels of quality. A channel is available in the highest level of quality if resources allow it. If resources were scarce, the channel would be offered in a degraded resolution.

D. Simulation parameters

The simulator takes parameters to perform a simulation run. We can set the number of users in their different categories.

Algorithm 1 Reduce Channel Slices

```
if activeChannelSlice >= Nslice then
    for i = 2 to totalLevel do
        for j = i to 2 do
            for all channels such that channel.level = j do
                if channel.mode ≠ DVR then
                    channel.level ← j − 1
                if activeChannelSlice < Nslice then
                    break
    if activeChannelSlice < Nslice then
        break
    for i = 2 to totalLevel do
        for j = i to 2 do
            for all channels such that channel.level = j do
                channel.level ← j − 1
                if activeChannelSlice < Nslice then
                    break
```

TABLE I
NUMBER OF CHANNEL SLICES ALLOCATED FOR DIFFERENT LEVELS OF QUALITY

Quality level	Number of channel slices allocated
1(lowest level)	1
2	2
3	3
4	4
5	5
6(4K level)	6

We can et the number of available content channels and the number of available slices. In addition, we can provide simulation time, zipf shape parameter and the means for channel flips and watch times.
Based on previous experimentation done in [1], we chose a zipf shape parameter of 0.9 for all experiments. We fixed the average watching time to 720 seconds (12 minutes), and the average channel flips to 6 channels, and the average wait time between flips to 2.

IV. EXPERIMENTATION

The objective of this research is to implement a simulation model with a multiple-level quality mode, to develop a dynamic algorithm to adjust the channel slices allocated for each content channel, dealing with the trade-off between efficient use of bandwidth and good viewing experience.

As for ”efficient use of bandwidth”, it is reflected by reducing the total channel slices, representing the total bandwidth requirements, as much as possible.

We measure ”good viewing experience” according to:

- Low blocking probability.
- High average playing quality experienced for all users.

More bandwidth implies higher playing quality, we use the number of channel slices allocated for a content channel to represent the playing quality of this channel. The way to measure the average watching quality for all users is to adding the average quality that an individual user experienced during the simulation and dividing the sum to get the average.

We would like to achieve high rewards to the ”stable” users and channels. Here ”stable” means that a channel is continuously being watched by a DVR user or a particular user does not flip too often. Meanwhile, the more ”popluar” channels, which is being viewed by more users should be allocated with more channel slices, representing higher quality. Therefore, the number of channel slices allocated for each content channel is adjusted according to the following 3 rules.

- A content channel, which is newly activated, enters into the lowest quality, allocated with one channel slice.
- A content channel will enter into a higher quality or a lower quality whenever the number of users watching it increases or decreases across a corresponding threshold (See Table II).
- A content channel, which is being viewed by a DVR user, enters into DVR mode. A content channel in DVR mode

TABLE II
THRESHOLDS BASED ON 6-LEVEL QUALITY STRUCTURE

Quality level	Threshold(percentage of users)
1(lowest level)	0.00%
2	3.125%
3	6.25%
4	9.375%
5	12.5%
6(4K level)	15.625%

has a corresponding DVR quality. Normally, a content channel in DVR mode will only increase its quality according to rule 2, but not dropping below DVR quality. In our system, a channel in DVR mode is allocated with at least 4 channel slices.

Currently the system has 6 levels of quality. The number of channel slices allocated for each level is showed in Table 1.

Based on rule 2, the threshold, which is the delamination between two adjacent levels, should be defined as a certain of number of users who are watching the specific channel concurrently. However, in the system, the threshold is implemented as a ratio, representing the percentage of total users, because the number of the whole users in the system could vary in multiple simulations. Table II shows the settings for thresholds based on the 6-level quality structure.

For example, the threshold between level 1 and level 2 is 3.125%. Then for a specific content channel temporarily in level 1, it will step into level 2 as soon as the number of users watching this channel exceeds 3.125 percents of the total users. On the other hand, the content channel will re-enter into level 1 from level 2 if the number of users watching it drops below 3.125 percents of the total users.

A. Runs without multiple quality levels

We want to demonstrate that multiple quality levels allows a service provider to mitigate the effect of allocating lower bandwidth. We ran two sets experiments with the same parameters: 75 users, 30 content channels, and we vary the amount of available channel slices in these two experiments. The only difference is that one experiment was run without multiple quality levels. As expected, we can see that allowing multiple quality levels reduces the amount of blocking in the two experiments as shown in Figure 3.

B. Runs with multiple quality levels

Since we want to avoid blocking, we concentrate the rest of our experiments on multiple quality levels. The obvious issue that we may run into with this approach is a reduced overall quality level. Therefore, we want to analyze quality levels while reducing blocking probability to less than 1%.

The overall goal is to give network operators an idea on how might they assign bandwidth to a given pool of users and content channels while minimizing blocking and achieving a reasonable level of video quality. Operators may have more users joining in, requiring an adjustment to the amount of bandwidth allocated. We ran multiple experiments with different bandwidth sets while varying the number of users to see how far can we reduce the bandwidth until performance is degraded to unacceptable levels.
We run a test with 100 channel slices, and 200 content channels, varying the number of users. This is an attempt to simulate a scenario where more users subscribe with a service providers. We notice that with 350 users, blocking is still about 1% or less. The quality level is close to 2, but has not degraded much from the scenario where we service 100 users as seen in Figure 5. Blocking and quality level are also affected by the number of available content channels, but providers do not increase the number of available channels very rapidly, and sometimes available channels remain fixed with minimal additions for years. Therefore, the number of users becomes a better indicator of how much a network operator should assign to meet their demands and improve the quality of experience.

We imagine that network operators are involved with setting bandwidth for a network given the number of channels and subscribers. Therefore, we expect the operator to find a good ratio of slices allocated to subscribed users and to offered content channels to be indicative of the performance of the system. We further analyzed our results to see if there is a correlation between slice-to-user ratio and slice-to-channel ratio to blocking and overall video quality. We picked test cases where blocking is lower than 1% and plotted the average quality levels against the user-to-slice ratio (see Figure 4), and concluded that there is a correlation as expected. For example, if we desire minimal blocking and a quality level of 3, we can allocate have a user-to-slice ratio of approximately 1.2; in other words, we can service 120 users with only 100 channel slices.

C. Effect of flipping behavior

We discussed earlier different categories of users. We focused on the effect of flipping behavior on the performance of the system. We ran a set of experiments with a fixed bandwidths and increased the number of users with each experiment to stretch the limits of the allocated bandwidth. The purpose is to see which user behavior exhibits a predictable

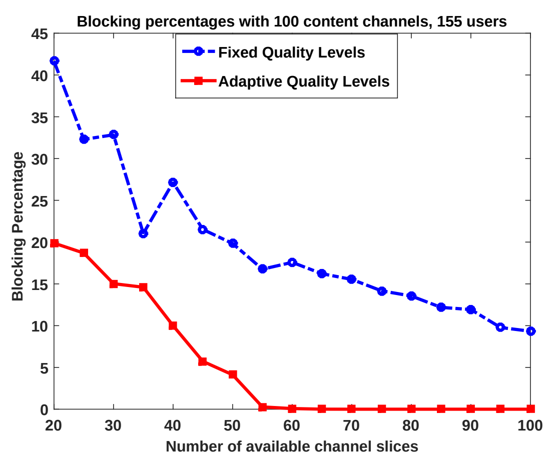

Fig. 3. Mitigating blocking by reducing quality

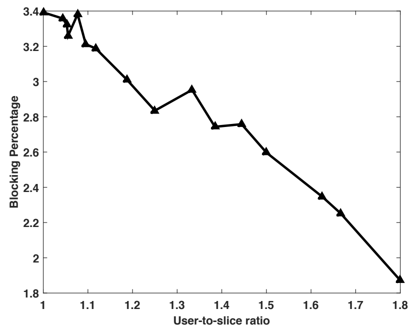

Fig. 4. Correlation between User-to-slice ratio and average video playback quality

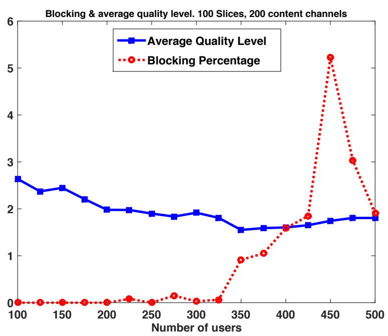

Fig. 5. Effect of users joining a system with limited bandwidth

and reasonable performance. We observe in Figure 6 that flipping randomly according to channel popularity generated from a Zipf distribution yeilds slightly lower blocking and is somewhat more predictable than flipping up and down through channels.

V. Discussion

We noticed from our experiments that users flipping randomly through channels that are ranked according to a Zipf law incurs lower blocking percentage that users that flip through channels in a sequential manner. As discussed in section II, research has been done to find an optimal ordering

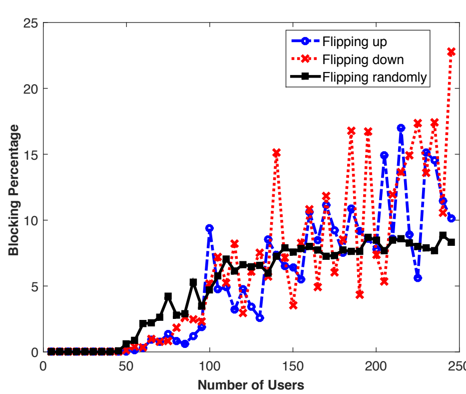

Fig. 6. Flipping randomly (according to Zipf ranking) shows more predictable blocking percentage over the long run

of channels. We recommend that service providers should invest in researching ways to improve the usability of their set-top boxes and remote control to make it easy and intuitive to add channels to a favorite list.

All our experiments were run with the assumption that 25% of the user's population are DVR users. This is because we have not found any reference or data about DVR use in a switched digital broadcast deployment. In addition, the results might be better when an operator has actual data collected from users of existing systems. Better results would insure that service providers can provide high level of service to popular high resolution channel, while at the same time still provide a reliable service to legacy users still user lower resolution channels.

VI. Conclusion and Future Work

We implemented a switched digital broadcast simulator and ran a multitude of experiments to study the benefits and the trade offs of using switched network infrastructure to facilitate video broadcast. We show that service providers can provide a reasonable quality of service level to subscribers with minimal blocking. There is a correlation between the ratio of allocated slice-to-user ratio and video quality given that we want a desired blocking percentage under 1%. The way to find this ratio might be elusive and require further research into the problem.

In the future we would like to profile our simulator to analyze the each channels load over a period of time and come up with adaptive thresholds for quality levels.

Acknowledgment

Clemson University is acknowledged for generous allotment of compute time on Palmetto cluster.

References

[1] J. Gong, D. Reed, T. Shaw, D. Vivanco, and J. Martin, "Qam resource allocation in mixed-format vod systems," *Journal of Communications*, vol. 2, no. 1, pp. 1–9, 2007.

[2] L. Breslau, P. Cao, L. Fan, G. Phillips, and S. Shenker, "Web caching and zipf-like distributions: Evidence and implications," in *INFOCOM'99. Eighteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, vol. 1. IEEE, 1999, pp. 126–134.

[3] J. Gong, D. A. Vivanco, and J. Martin, "Simulation study of the spectral capacity requirements of switched digital broadcast," in *Broadband Communications, Networks and Systems, 2007. BROADNETS 2007. Fourth International Conference on*. IEEE, 2007, pp. 405–414.

[4] A. Taylor and R. Harper, "Switching on to switch off: An analysis of routine tv watching habits and their implications for electronic program guide design," 2001.

[5] T. Wauters, J. De Bruyne, L. Martens, D. Colle, B. Dhoedt, P. Demeester, and K. Haelvoet, "Hfc access network design for switched broadcast tv services," *Broadcasting, IEEE Transactions on*, vol. 53, no. 2, pp. 588–594, 2007.

[6] S. Joe Matarese, "New audience share opportunities: Advanced advertising, network dvr, and switched digital."

[7] L. Natali, C. Barcellona, and M. Merani, "Leveraging users' likes in a video streaming p2p platform," in *Computer Communications Workshops (INFOCOM WKSHPS), 2014 IEEE Conference on*. IEEE, 2014, pp. 269–274.

[8] T. Qiu, Z. Ge, S. Lee, J. Wang, J. Xu, and Q. Zhao, "Modeling user activities in a large iptv system," in *Proceedings of the 9th ACM SIGCOMM conference on Internet measurement conference*. ACM, 2009, pp. 430–441.

[9] N. Sinha, R. Oz, and S. Vasudevan, "The statistics of switched broadcast," in *SCTE 2005 Conference on Emerging Technologies, Tampa, USA*, 2005.

[10] J. Gong, D. Reed, T. Shaw, D. Vivanco, and J. Martin, "Vod qam resource allocation algorithms," in *NETWORKING 2006. Networking Technologies, Services, and Protocols; Performance of Computer and Communication Networks; Mobile and Wireless Communications Systems*. Springer, 2006, pp. 268–280.

[11] H. Yu, D. Zheng, B. Y. Zhao, and W. Zheng, "Understanding user behavior in large-scale video-on-demand systems," in *ACM SIGOPS Operating Systems Review*, vol. 40, no. 4. ACM, 2006, pp. 333–344.

[12] J. Gunter, "A future switched video system," *LCS, IEEE*, vol. 1, no. 1, pp. 66–72, 1990.

[13] H. Joo, H. Song, D.-B. Lee, and I. Lee, "An effective iptv channel control algorithm considering channel zapping time and network utilization," *Broadcasting, IEEE Transactions on*, vol. 54, no. 2, pp. 208–216, 2008.

[14] F. M. Ramos, J. Crowcroft, R. J. Gibbens, P. Rodriguez, and I. H. White, "Reducing channel change delay in iptv by predictive pre-joining of tv channels," *Signal Processing: Image Communication*, vol. 26, no. 7, pp. 400–412, 2011.

[15] H. Kim, J. Crowcroft, and F. M. Ramos, "Efficient channel selection using hierarchical clustering," in *World of Wireless, Mobile and Multimedia Networks (WoWMoM), 2012 IEEE International Symposium on a*. IEEE, 2012, pp. 1–9.

[16] Z. G. Kingsley, "Selective studies and the principle of relative frequency in language," 1932.