Under the Plastic: A Quantitative Look at DVD Video Encoding and Its Impact on Video Modeling

Wu-chi Feng, Jin Choi, Wu-chang Feng, Jonathan Walpole
Dept. of Comp. Sci. and Engineering
OGI School of Science and Engineering at OHSU
20000 NW Walker Rd.
Beaverton, OR 97006
{wuchi,jchoi,wuchang,walpole}@cse.ogi.edu

In this paper, we examine the DVD encoding process and the implications this process has video modeling and network traffic analysis. We have assembled a system that allows us to extract the video data from the DVDs as they were encoded for distribution. Analyzing the resulting video trace data, we describe how DVD encodings have evolved over time. In addition, our findings show that the underlying video content is fundamentally different than those produced by basic consumer video capture boards. We demonstrate how this affects current video modeling proposals and their affect on network traffic characterization. This research is based upon the analysis of over 100 DVD video streams, whose statistics will be made available after the Packet Video workshop[1].

Keywords
DVD encoding, MPEG-2, video content

1. INTRODUCTION

As computing and networking technologies continue to advance, the ability to stream high-quality video data across networks is becoming increasingly viable. As a result, it is envisioned that delivering DVD-quality video streams directly from the content provider will soon become a reality, assuming the challenges of digital media copyright protection can be effectively solved. In order to more fully understand the implications of this on current networking, operating systems, and multimedia file systems design, this paper attempts to more fully characterize the actual content being generated by the content provider. While the DVD encoding process has limitations of maximum bit-rate and the need to fit all the video data with maximal quality on a fixed-size platter, we can easily envision content providers using the DVD source content for distribution to limit re-encoding effort and extra man-hour time. More specifically, we examine the DVD encoding process itself in an attempt to understand how content is evolving and how solutions that are being developed at conferences like Packet Video will be affected by the DVD creation process.

While the characterization of video sources has been widely studied in the literature, most of these studies have taken content from VHS tapes, laser discs, and DVDs and have re-encoded them using an analog video capture device. In addition, almost all of these have employed fixed patterns of video frame types generated by either software-only encoders or hardware-based encoding boards. As our analysis of the DVD data shows, this comes far from capturing the reality of modern content libraries.

In this paper, we present our observations of compressed MPEG-2 data streams taken directly from consumer DVDs. We have implemented a system to extract MPEG-2 data from DVD content in its native format. Using the data from nearly 80 DVDs, we have compiled a large library of MPEG-2 statistics for analysis. This library consists of 107 video streams representing approximately 140 hours of video. These sequences range from half-hour television shows converted to DVD to full-length feature films. Using this data, we highlight several trends that we observed in the DVD encoding process. We then analyze the burstiness and self-similarity of these streams and its impact on network traffic generation and video modeling. Our results indicate that much more work needs to be done in video modeling or that researchers ought to use real trace data for their research studies.

In the next section, we will describe the basic mechanisms that we used to capture the video data and how this differs fundamentally from analog capture that has been the source of many video traces used in previous studies. Section 3 summarizes and describes the video data captured. Section 4 describes some of the trends that we observed through analyzing the video streams. In particular, a discussion of some of the ramifications to network and operating systems research is discussed. Finally, a conclusion and directions for future work are provided.

[1] If the program committee thinks the availability of the data at the workshop would be useful, we can have CD's made of the statistics before the conference and have them available for discussions

Contributions of this work: We believe that there are several main contributions of this work. First, we have performed an in-depth analysis of DVD video content and its impact on network services. Second, we show how video models need to evolve in order to capture the essence of DVD video encoding. Third, we will contribute a significant library of trace data (both DVD-based and software-encoding-based) to the community for further research using real video trace data.

2. BACKGROUND

2.1 Capturing Video Data

There are a number of ways to create a library of compressed digital video streams. To create the highest quality streams, having the original source film or digital content provides both the best quality video for compression and allows for meaningful comparisons like peak signal-to-noise ratio measures. Obviously, the tight grip of the content creators makes this nearly impossible to create a large library of such streams. Another way to create such a library is to take analog video data from devices such as a DVD player, a laser disc player, or a VHS tape player and then capture and compress the stream. For the compression into the target video format, one can then use a hardware compression board or software to create the final MPEG format video. This approach, however, suffers from a number of limiting factors. First, film-based movies are typically shown at a rate of 24 frames per second. When converted to an analog signal suitable for capture, it is typically converted (for an NTSC signal) to 60 fields per second, significantly altering the characteristics of the content (or making the conversion process incredibly slow). Second, and more importantly, the conversion into analog can cause interlacing artifacts especially when the compressed content is of higher quality (such as the DVDs 720x480 pixel resolution). While these artifacts generally do not appear in the smaller-sized MPEG-1 encodings that have appeared in the literature (such as [3, 7]), they present a substantial hurdle in compressing higher quality streams. Finally, these techniques typically use fixed frame type patterns for the entire sequence. As we will show, this is hardly the case for the data found on modern DVDs.

For our work, we have chosen to use the original DVD content itself and to stay completely within the digital domain. By doing so, we can directly access the digital data without the frame rate conversion problems and without the interlacing artifacts. The data captured also gives a more realistic picture of the content that content providers will eventually provide over the network. Specifically, the compressed MPEG-2 video we capture is in a format tailored to the content creator's specification[2] and has not been modified. We will describe this video data in the following sections.

The system we implemented consists of two main components: a ripper which retrieves the video data off of the DVD and an analyzer which decompresses the video data and saves statistics regarding the video data. We used the program *SmartRipper* to extract the information (IFO) files and the video object (VOB) files from the DVD. We then ran the VOBs through a modified version of the *Flask* MPEG decoder which only parsed the video data and skipped the computationally expensive inverse DCT and reconstruction operations. For each frame of video, we then saved the frame type, the frame size, and macroblock information. For the latter, we saved the macroblock type, the quantization value used, and the presence of forward and backward motion vectors. This resulted in approximately 400 megabytes of data per two-hour movie for the macroblock information (in binary format) and approximately 2 megabytes for the frame size and frame type information (in ASCII format). Using this process, we were able to create the statistics files for each video in approximately 20-25 minutes, split fairly evenly between ripping and decoding.

In addition to the streams captured through the method above, we have also taken a number of the streams and re-encoded them using the *Flask* software MPEG-2 encoder to compare and contrast constant quality encoding of video with the video found on the DVDs. We will describe this in further detail in the experimentation section.

3. THE VIDEO DATA

Our library consists of 76 DVD's that make up approximately 140 hours of video. The DVD data was then separated into their individual viewing units resulting in 107 video sequences. The increase in sequences comes from separating collections of smaller TV shows combined onto one DVD. For example, each *Best of Friends* DVD has five half hour video clips. In all, the total data collected represents over 300 Gigabytes of compressed MPEG-2 video. The DVDs cover a large classification of titles including:

- Children's titles - including DVDs such as various *Barney* movies, *Dumbo*, *Sesame Street - 25th Anniversary Celebration*, *Elmo in Grouchland*, and the *Tigger Movie*.
- Made from television titles - including the *Best of Friends*, the *Simpson's First Season*, *The NBA's 100 Greatest Plays*, and *South Park*.
- Computer generated movies - including *Bug's Life*, *Book of Pooh*, *Shrek*, *Toy Story*, and *Toy Story 2*.
- Action movies - including movies such as the *Die Hard* series, *Independence Day*, *Bad Boys*, and *Mission Impossible*.
- Dramas - including films such as *The English Patient*, *Hunt for Red October*, *You've Got Mail*, and *The Godfather (Parts I - III)*.
- Older films - including *Mary Poppins*, *Fantasia*, *Princess Bride*, and *Caddy Shack*.
- Music Video DVDs - including *James Taylor*, the *Indigo Girls*, and *Phish at Beacon Hill*.

A complete listing of the video sequences is shown in Figure 9 at the end of paper. Finally, we note that because of copyright legalities, the collection of DVDs presented in this paper represents our own personal DVD collections. We are in the process of adding several more video streams to

[2]With the DVD storage capacity and maximum bit-rate of approximately 10 Mbps in mind

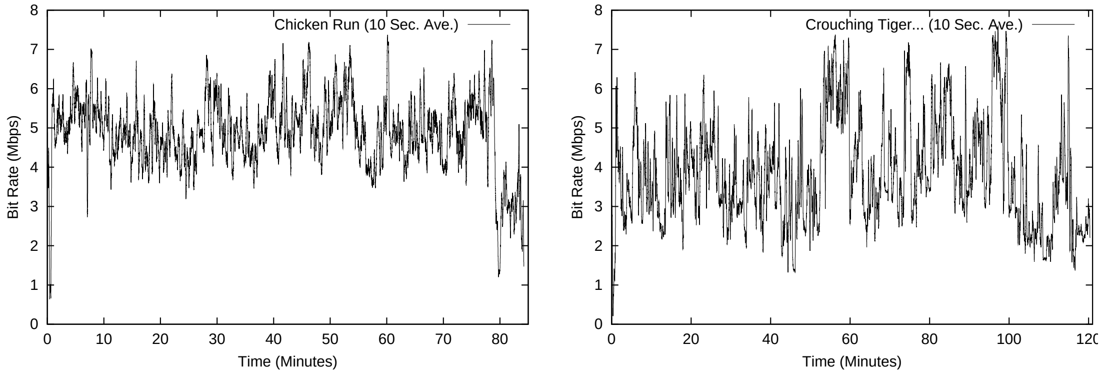

Figure 1: Video Data Examples. This figure shows two example video streams that were obtained from the DVDs. The left figure shows the bandwidth requirements for the movie *Chicken Run* while the right figure shows the same for the movie *Crouching Tiger, Hidden Dragon.*

round out the collection but believe that we have a fairly representative set of video as will be shown in the experimentation section). An example of the video data captured is shown in Figure 1. As shown in this figure, there is typically a large amount of variation in bandwidth requirements for the DVDs over time.

4. TRENDS AND DISCUSSION

We begin our discussion with general observations with the encoding of DVD video content in Section 4.1. We will then focus on a more in-depth analysis of the structure of the video data, an analysis of the burstiness and self-similarity of the video data, and a comparison of DVD video content compared to software-encoded video data in the rest of the section.

4.1 General Observations

4.1.1 Bit-rates

The DVD format is limited to approximately 4.7 Gigabytes per side per layer. Some of the DVDs were encoding using dual layer format, while others packed the entire video stream into a single layer. Which ever layer format was chosen, the video data was typically encoded to fill the entire DVD up with data. The range of average bit-rate values we saw were from 3.3 Megabits per second to 7.8 Megabits per second. Because the encoding of the video was more or less fixed by the size of the DVD, this seems to indicate that given a larger DVD capacity, that even larger bit-rates would be desirable for the streams. We also note that the bandwidth limit of the DVD format is approximately 10 Mbps (determined by the spin speed of the DVD).

4.1.2 Sequence and GOP headers

We observed that each group of pictures (GOP) had only a single I frame in it for all the videos. An interesting observation is that, without exception, each of these GOPs was encoded into a separate sequence. That is, each I-frame is immediately preceded by a GOP header, which in turn, is immediately preceded by a sequence header. We believe that these might be used to support efficient fast-forward and fast-rewind operations on the DVD players. Finally, the GOP sizes for all movies tended be around 12 frames long, or half the original frame rate. For TV content the GOP size was typically 15 frames in length. As we will describe shortly in Section 4.2, the number of frames in each GOP was originally fixed but has become much more adaptive over time.

4.1.3 DVD encoding companies

For the most part, many of the DVDs do not identify the DVD encoding company used in the encoding process. Fourteen of the DVDs, however, have a separate chapter at the end of the DVD that identifies the encoding company. The three most prevalent were the Digital Video Compression Center, Laser Pacific Media Corporation, and Panasonic Disc Services Corporation. While we expected that these companies would be associated with a particular movie company, this was not the case. For example, DVCC has compressed data for 20th Century Fox, Columbia, and Universal Studios.

4.1.4 CBR vs. VBR video coding

As one would expect, nearly all the bandwidth requirements of the video streams were quite variable over time. Our observations indicate that while most of the streams have significant bandwidth variability over time, they are not encoded with constant quality. That is, the quantization values change over time. Presumably, we believe that the quality is adapted over time with respect to some peak signal-to-noise ratio measurement. The notable exception to this was the *Spy Kids* movie and some of the TV show encodings. For these particular DVDs, the bandwidth requirements were essentially constant-bit-rate for the entire DVD. We also note that the *Spy Kids* movie is the only movie in the dataset from the movie company Dimension, which seems to be relatively new to the movie industry. In addition, our observations lead us to believe that they used a relatively non-standard DVD encoding company to make the DVD.

We are still analyzing the data to understand how these parameters are chosen.

4.2 Frame Patterns & Group of Pictures

The frame patterns and group of pictures grouping were originally designed to allow encoders to have a great deal of flexibility in their encoding process. Many of the MPEG video streams that have been generated in the literature [3, 7, 2], and used by many systems and networking researchers have fixed group of pictures patterns for the entire length of the video data. In fact, some techniques have been proposed that take advantage of fixed frame patterns within MPEG-compressed video streams (e.g. [4]).

For the DVD collection assembled here, we have observed that both the number of frames within a group of pictures and the frame patterns within the stream can vary quite a bit. In fact, newer movies tend have much larger variations in GOP patterns (more on this later). As an example of how much the patterns and GOP sizes change, the movie *Independence Day* has 134 unique GOP patterns. Some of the more interesting ones from this movie include:

- *IBBPBPBBPB* (occurred 134 times)
- *IBBPPBBPBBPBB* (occurred 194 times)
- *IBPPPBPPBPPB* (occurred 64 times)

In all, we observed 526 unique GOP patterns across all movies. This is quite surprising given that nearly all the encoded video streams have GOP sizes that are around one half second long. For movies encoded at 24 frames per second, the most common sequence of frame patterns was $IBBPBBPBBPBB$ without exception. We suspect that the actual choice of P or B frame type used is based upon the result of the motion estimation searching, a peak signal-to-noise ratio measure, or a scene-detection algorithm.

In terms of the use of GOPs, we found that each video stream exhibited one of six characteristic GOP length patterns over time. In the simplest case (Figure 2(a)), the video was encoded using a fixed GOP pattern for the duration of the video. This encoding occurred only with the *Friends* and *Simpsons* videos. The second type of GOP length pattern characteristic (Figure 2(b)) observed consisted of GOPs of mostly the same length, but with some of the GOPs truncated. In these videos, the number of truncated GOPs was highly correlated with the number of chapters in the video. That is, they only occurred at the end of chapters. Note that when the chapter and GOP pattern were aligned, there was no shortened GOP. In the third type, the GOP pattern observed was similar to the second but allowed for GOPs which extended beyond the typical 12 frames. The fourth type of GOP length pattern (Figure 2(d)) had only truncated GOPs but at much higher frequency than at the chapter level. This pattern appeared to be the first stage in which adaptive GOP patterns were employed as basic scene detection algorithms began to appear. The fifth type of GOP had a fairly strict cut-off on how large or small a GOP could be. In the example shown in Figure 2(e), a GOP could either be extended by 2 frames or truncated to create a GOP of at least 4 frames. Finally, the sixth type of GOP length pattern had no observable limitations and exhibited a highly variable GOP length and frame pattern throughout the stream. For each video examined, the table lists the GOP encoding characteristic which has been observed for the video. We will use the GOP type classification in the latter sections of this paper.

The GOP length encodings show the evolution of scene detection algorithms over time. For a single movie company like Columbia Pictures, the pre-2000 *release* DVDs tend to have type 2 encodings; the 1999-2000 *release* DVDs tend to have either type 4 or type 5 encodings; and the 2001 and on DVDs have type 6 encodings. While we expected that the newer DVDs would all have GOP length pattern type 6 encodings, we found that the GOP pattern encoding was somewhat tied to the movie company. For example, Columbia and 20th Century Fox use completely adaptive GOP lengths in all of their recent movies, however, the most recent movies by release date (*Shrek*, *Stargate*, and *Rush Hour 2*) used relatively simple GOP encodings. Note these movies also come from companies that are relatively new, indicating that they may be creating their own DVD content and are going to evolve to the more advanced encodings with each new release.

4.3 Scene Detection Algorithms

Scene detection algorithms have been proposed for video storage and retrieval systems. When encoding video data, if the beginning of each scene can be detected, then a new GOP can be started for this scene. Because GOPs are the random access mechanism within MPEG video streams, breaking the video into the right scenes helps coding efficiency as well as random access.

From our observations, scene detection algorithms are making their way into the DVD encoding process. Having examined over two hours of video from Mulan and Mission Impossible, we found that nearly all of the scenes that we examined started correctly with an I-frame. In addition, in the newer encoding that use adaptive GOP patterns and lengths, we found that GOPs with a single I-frame followed by a sequence of only P frames (e.g. $IPPPPPPPP$) are quite common. These tend to occur in the following situations: (i) when the scene is extremely short in length (e.g. a quick scene change in a fight), (ii) fade in and fade outs between scenes, and (iii) scenes with explosions in them (e.g. thunder). Thus, our observations indicate that the basic scene detection algorithms work quite well for the DVD encoding process.

4.4 Effect of DVD Coding on Network Resources

What we discover on DVDs in terms of the diversity of encoding leads to an interesting question: What impact does this have on video modeling and on network requirements for streaming DVD content over the Internet? At first glance, it would make sense that the type of GOP pattern used will directly influence the burstiness of the video stream. Previous studies characterizing VBR video have either fixed the GOP pattern being used or augmented it with simple scene-based variations in structure. With the emerging encoding techniques being applied to VBR video, we first look at how such encoding is impacting the variability of video. In order to measure variability across multiple time scales, we use the standard aggregated variance method to estimate the Hurst parameter (H) of the frame size data from the DVDs [6]. In this method, the sequence is divided into multiple, consecutive, equally-sized, blocks. The values within the block are averaged and the variance of the sequence of

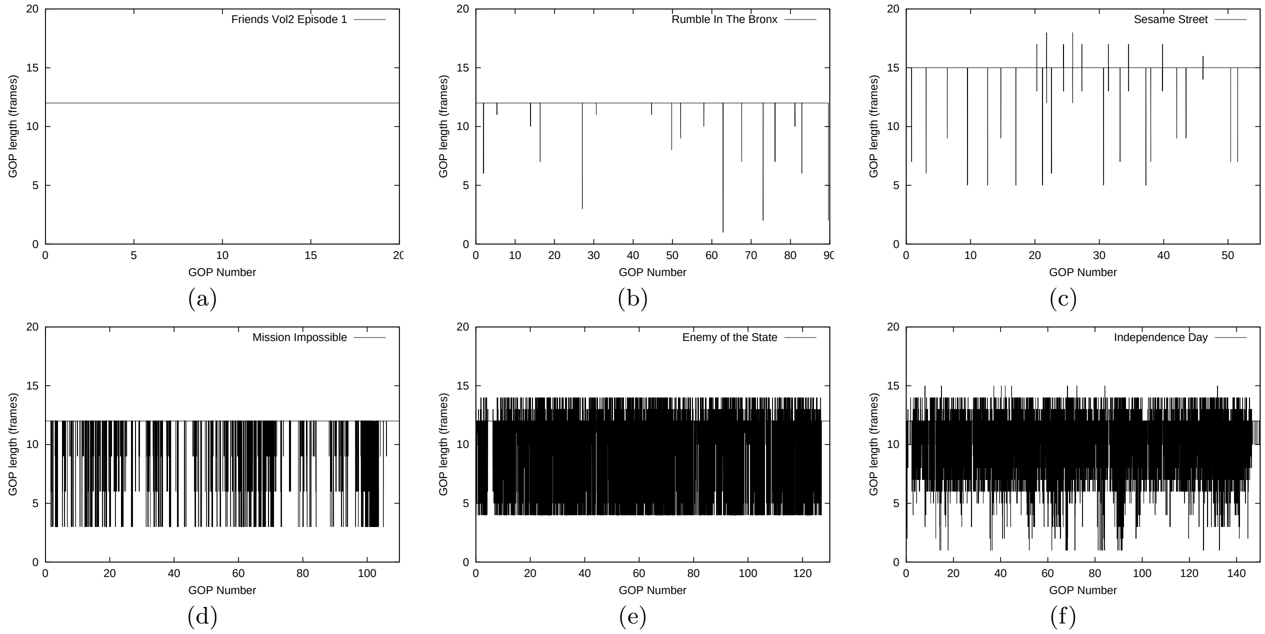

Figure 2: GOP Length Versus Time Plots.

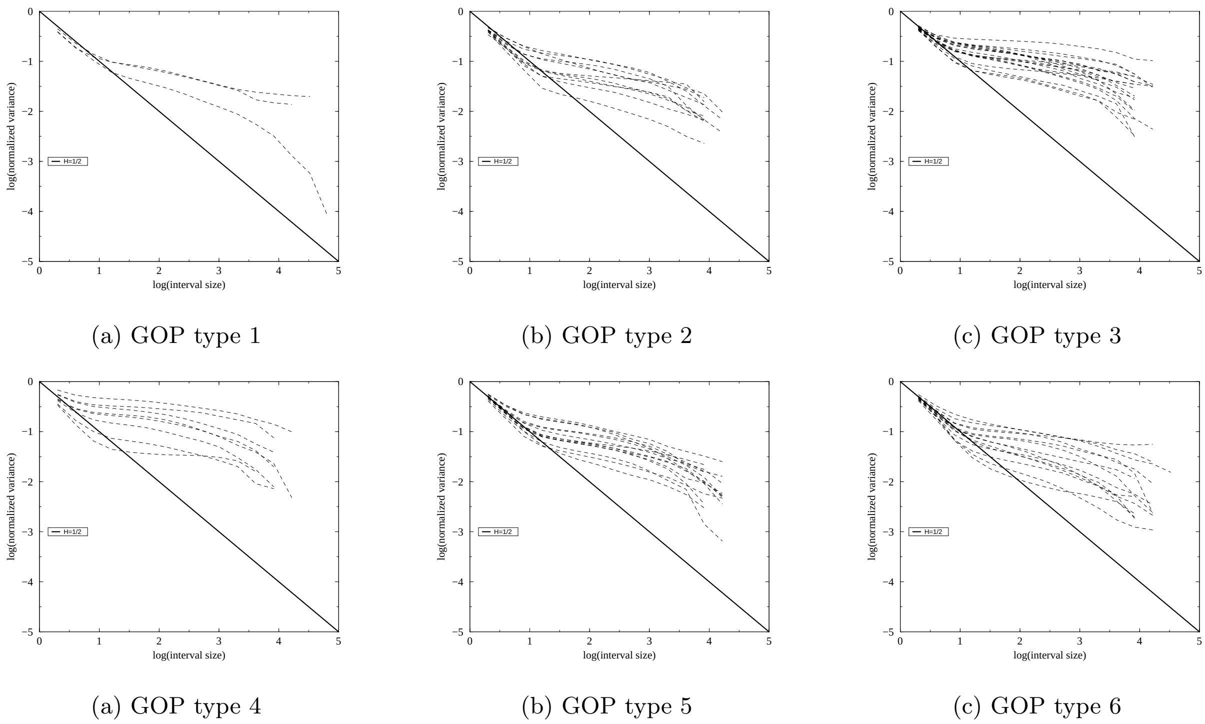

Figure 3: Aggregate variance plots for each DVD set based on type

averages is calculated. For a short-range dependent process, as the block size is increased, the variance of the resulting sequence consistently decreases. In contrast, for long-range dependent sequences, the sequence maintains high variability across block sizes and time scales. To determine the degree of long-range dependence, the log of the normalized variance is plotted against the log of the block size. The normalized variance is calculated as the variance of the aggregated sequence divided by the variance of the initial, unaggregated sequence. The block size, in this case, is the number of frames per block. The Hurst parameter (H) can be estimated by taking the magnitude of the slope of the best-fit line through the data points (β) and calculating H via the relation $H = 1 - \frac{\beta}{2}$. The Hurst parameter thus ranges between $\frac{1}{2}$ and 1. A short-range or no dependence sequence will have a slope of -1 which corresponds to a Hurst parameter of $\frac{1}{2}$ while a long-range dependent sequence will have a value closer to 1.

Figure 3 shows the aggregate variance plots of the movies in each of the 6 GOP encoding classes. To our surprise, it appears that the GOP encoding scheme may not, in fact, have any impact on the video variability of the DVDs in our archive. The figure shows that while the encoding structure may vary, the frame sequence variability remains unpredictable with no direct correlation between encoding type and burstiness. This is key as structurally it appears that GOP patterns have minimal influence on overall frame variability over large time scales.

To further illustrate this, we take a closer look at the *South Park* DVDs. The *South Park* DVDs are unique in that the TV volumes are encoded using a structurally fixed scheme while the movie is encoded in a structurally variable scheme. (For our treatment of the *South Park* TV volumes, we aggregate all of the separate volumes into one sequence.) Thus, while the content is relatively the same, the encoding algorithm used is completely different. Figure 4(a) shows this difference by plotting the GOP length encoding patterns of both DVDs. Figure 4(b) shows the variance plot of both movies. Note that while structurally, the encoding is extremely different, the movies exhibit similarly smooth variance behavior.

In contrast to the behavior seen in Figure 4, Figure 5(a) plots the aggregate variance behavior of the first two *Godfather* movies. Both movies are produced by the same studio, have similar content, and similar encoding structure types. As the figure shows, even with all things being relatively similar, the variance behavior differs dramatically with the first *Godfather* movie being significantly more variable than the second one over all time scales. Figures 5(b) and (c) show the GOP length encodings of the first 4000 frames of the *Godfather* and the *Godfather II* respectively. As the figure shows, both movies use a highly variable encoding structure. Finally, Figure 6 illustrates the average frame sizes of 32-frame blocks for the first 8000 blocks of each movie. Again, as corroborated by the aggregate variance plot, the original *Godfather* movie exhibits far more burstiness than the sequel.

4.5 A Comparison of DVD and Software-Based Encoding

To examine the differences between software-encoding and the DVD video content, we took three of the movies (*Spy Kids, Shrek, and Sneakers*) and re-encoded them using the *Flask* software MPEG video encoder. The purpose of this was to study the difference in using a constant quality and constant frame pattern in the software encoder versus the DVD content. To encode the video, we chose the frame pattern that occurred most often within the video stream in the DVD format. In this case it was $IBBPBBPBBPBB$. We also chose a fixed quantization level of 4 for these experiments. The movies were chosen to represent a CBR-like stream *Spy Kids*, an extremely bursty stream (*Sneakers*, and one in the middle *Shrek*.

The resulting re-encodings are shown in Figure 7. We have shown the bit rates in figures (a) - (c) for both the original (bold lines) and the re-encoded version (dotted lines), and have graphed the quality per macroblock of the DVD encoded stream for comparison in figures (d) - (f). The quality per macroblock is simply the average quality level used per frame considering only the encoded macroblocks. As expected, the movie *Spy Kids* (shown in figure (a)) is significantly different using a constant quality. The DVD version has a 6.98 Mbps average bit rate versus 5.97 Mbps for the re-encoded version. In (b) we have shown the movie *Shrek*. Here, we see that the DVD stream is similar to that of the constant quality stream (with more attenuation of the bit rate). We also note that the quality is also less variable over time than the *Spy Kids* video and that the average bit rate dropped from 7.06 Mbps to 4.97 Mbps in the re-encoding. Finally, the *Sneakers* video is shown in figure (c). We see that the quality of the stream is fairly constant and that the resulting video stream is similar to that of the DVD encoding. The average bit rate dropped from 7.06 Mbps to 4.97 Mbps for the constant quality encoded version.

One interesting note is that the end credits of the videos and their re-encodings were quite different. In the DVD encded movies, the end credits typically have a smaller bit-rate than the rest of the movie. In contrast, however, the re-encoded constant-quality streams have a bit-rate for the end credits that increases significantly. We somewhat expect this due to the typically high amount of high frequency data required to represent the text of the end credits.

To highlight where the difference in bit rate requirements of the re-encodings, we have graphed histograms of the re-encoded movies in Figure 8. The key difference between the DVD content (figures (a) - (c)) and the re-encoded content (figure (d) - (f)) is clearly in the distribution of the frames. For the DVD content, we see that for the DVD content of the constant bit-rate *Spy Kids*, the distribution of frames is heavily clustered for the B and P frames. As we move towards more constant quality encodings in the *Shrek* and *Sneakers* videos, we see that the frame distributions for the I, P, and B frames becomes more intermixed but still present. To contrast this, we see that for all of the re-encoded, constant quality streams that the distribution is more uniform. This suggests that in the modeling of DVD video data using multiple distributions (as in [5, 1, 8]) is necessary.

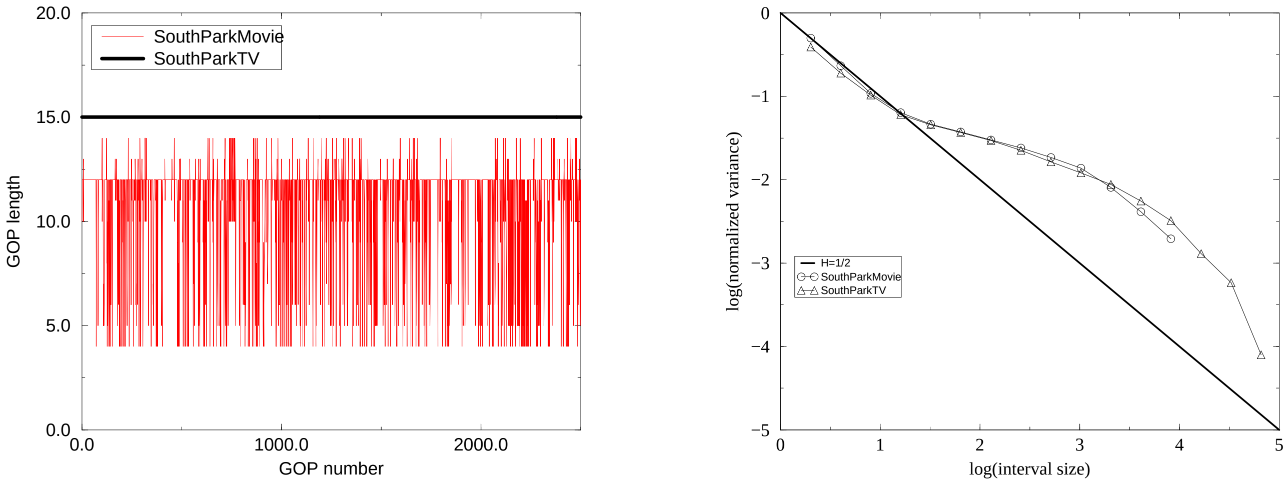

Figure 4: GOP encoding and aggregate variance plots for *South Park*

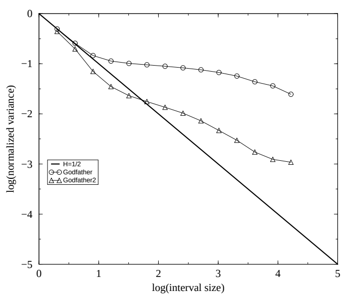

(a) Aggregate variance plot

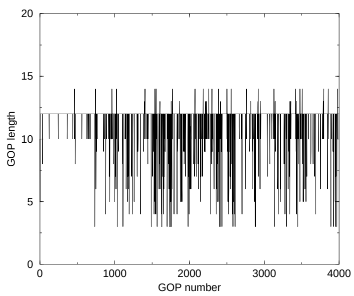

(b) GOP encoding for the Godfather

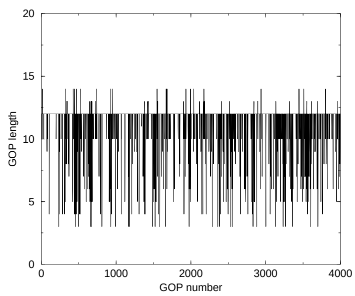

(c) GOP encoding for the Godfather II

Figure 5: Variance behavior and GOP encodings for the Godfather movies

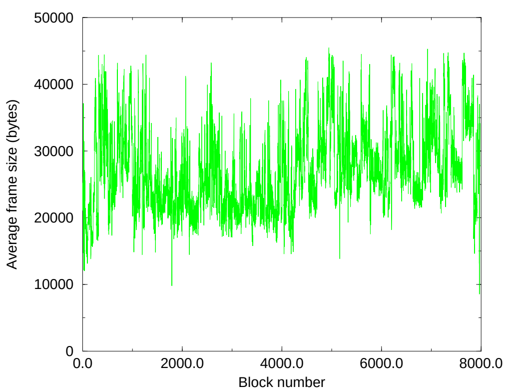

(a) Godfather

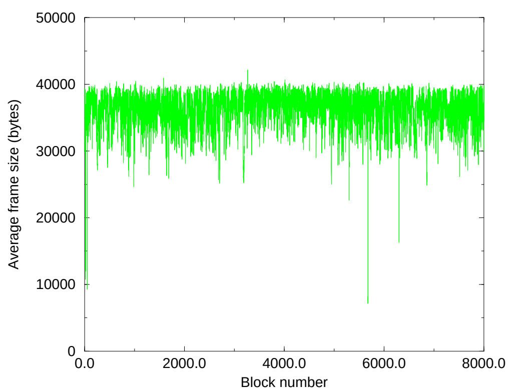

(b) Godfather II

Figure 6: Average frame size plot for 32-frame blocks of the Godfather movies

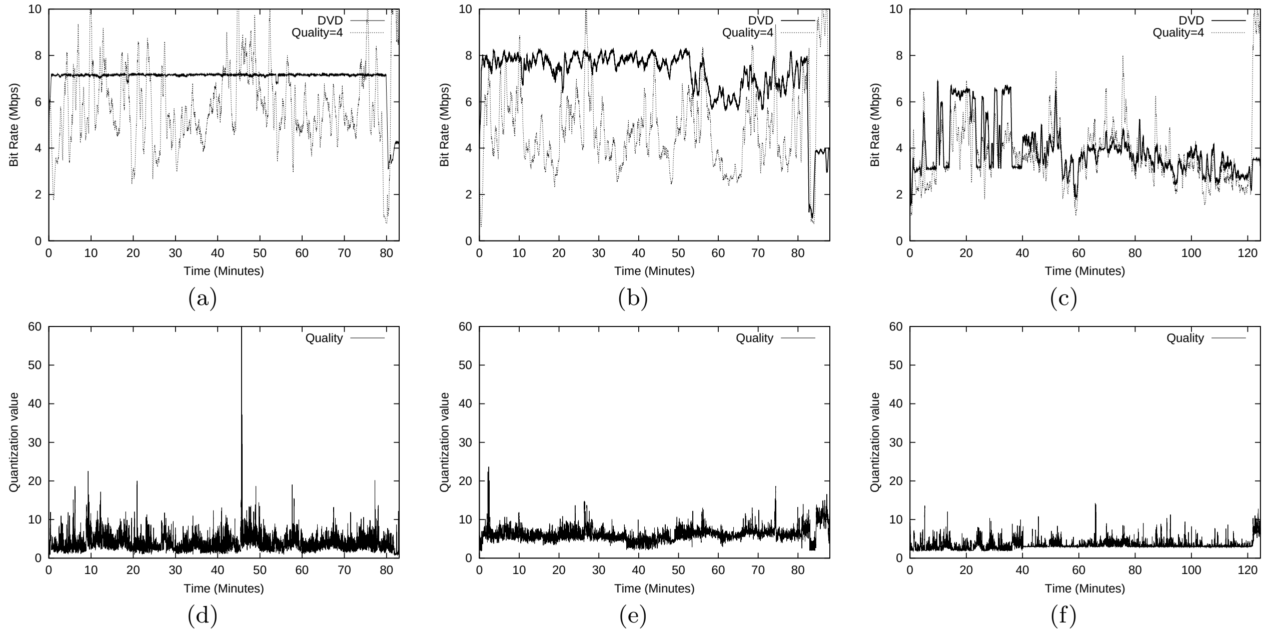

Figure 7: Video Re-encoding Examples. In (a) - (c), we have graphed the bandwidth requirements for the movies *Spy Kids*, *Shrek*, and *Sneakers*, respectively. In (d) - (f) we have graphed the average quality per encoded macroblock for the DVD encoded video content

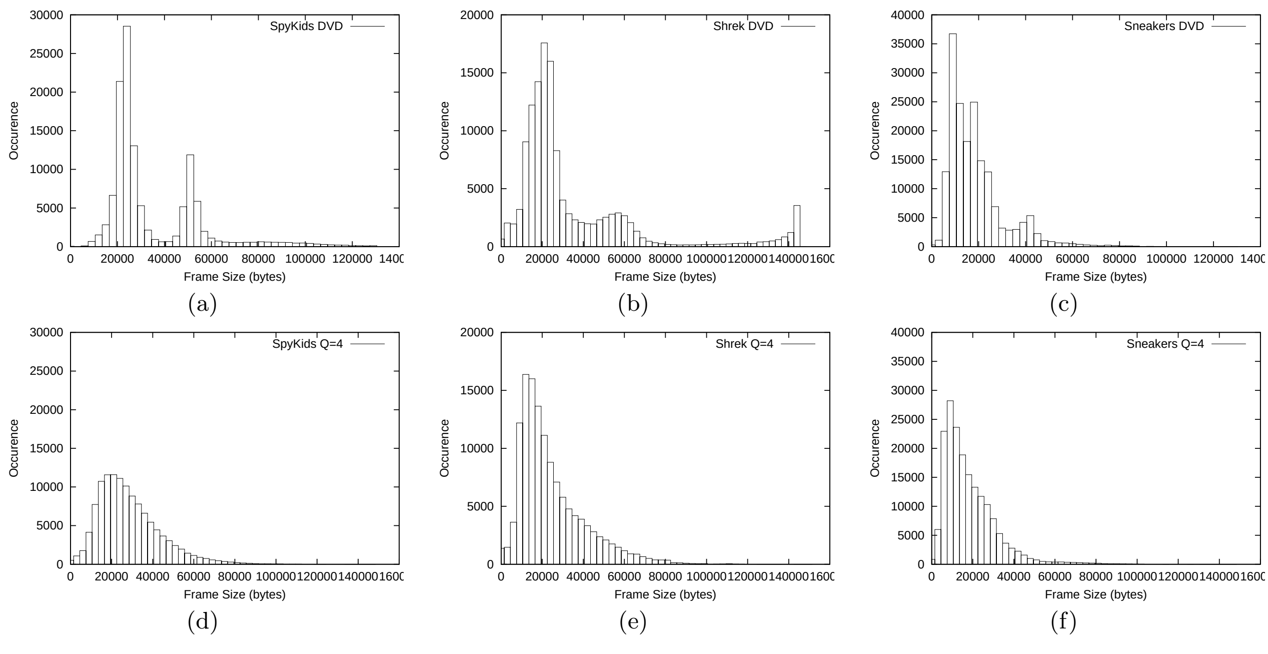

Figure 8: Distribution of Frame Sizes. Figures (a) - (c) show the frame distributions for the DVD encodings of the movies *Spy Kids*, *Shrek*, and *Sneakers*, respectively. In (d) - (f), we graphed the frame distributions for the re-encoded versions using a constant quality of 4.

4.6 Effect of DVD Coding on Systems Design

In this subsection, we briefly highlight some of the other implications of using DVD video content in multimedia systems and urge researchers to re-visit their research findings based upon the movie data represented here.

First, there are some interesting ramifications of using DVD video for streaming across networks. The most important of which is the presence of GOPs with significant runs of P-frames in the more recent DVDs. In many adaptation algorithms that have appeared in the literature, frame dropping is accomplished within a GOP by first dropping the B-frames and then the P-frames. Because many of the videos that we have seen in the literature follow a typical pattern such as:

$$IBBPBBPBBPBBIBBP$$

dropping B-frames can be done in such a way as to evenly distribute the removal of frames. For the long runs of P-frames that we are seeing in the DVD content, dropping of the frames within these GOPs must be done from the end of the GOP working sequentially backward in time, resulting in significant degradation in smoothness of playback. Thus, we expect that even the simplest adaptation algorithms that remove frames will need to deal with some form of transcoding if DVD content forms the basis of video libraries in the future.

In terms of operating systems support for DVD encoded video data, there have been some research efforts that have focused on the processing requirements to decompress MPEG and MPEG-2 video streams in software. One of the key observations is that the different frame types require differing amounts of processing time, as one would expect[9]. Having more or less random GOP patterns in the newer encodings has ramifications on the allocation of processor, memory, network, and I/O resources for DVD data. In particular, with the variability of frame patterns, one would expect that the processor resource will have to be less aggressively allocated to accommodate these patterns.

5. ACKNOWLEDGEMENTS

This research was funded in parts by National Science Foundation grants NSF ANIR 98-75493 and NSF EIA 01-30344. The results presented in this paper do not necessarily reflect the views of the funding agency. The authors also wish to thank Larry Rowe and Barry Haskell for their valuable input in the development of this paper.

6. CONCLUSION

In this paper, we have described a quantitative analysis of a large amount of DVD content in its original compressed form. Our observations indicate that the encoding process for DVD content is significantly different than the encodings that have appeared in the literature. The encoding of this data has implications on the design and implementation of network and operating systems that need to support the flow of these streams. Our results show that DVD content is evolving and becoming more complex in its encoding structure of frame types and frame patterns.

The work in this area has just begun. We need to understand how the DVD encodings affect video models that currently exist and how they should be extended to incorporate some rate control in the algorithms. We also need to understand how the adaptation process of video data over best-effort networks will be accomplished given the non-regular encoding sequences of frame types that we have observed. Finally, we will make the video statistics data available to the research community in general to start the work on understanding the impact of real trace data in the design of future systems.

7. REFERENCES

[1] K. Chandra and A. R. Reibman. Modeling one- and two-layer variable bit rate video. *IEEE/ACM Transactions on Networking*, 7(3):398–413, 1999.

[2] W. Feng. Pseudo-modeling of mpeg-based variable-bit-rate video. In *Proceedings of SPIE/ACM Multimedia Computing and Networking 2000*, January 2000.

[3] M. Garrett and W. Willinger. Analysis, modeling and generation of self-similar vbr video traffic. In *Proceedings of SIGCOMM 1994*, pages 269–280, September 1994.

[4] M. Krunz and S. K. Tripathi. Exploiting the temporal structure of MPEG video for the reduction of bandwidth requirements. In *Proceedings of INFOCOM 1997*, pages 67–74, 1997.

[5] M. Krunz and S. K. Tripathi. On the characterization of VBR MPEG streams. In *Proceedings of SigMetrics 1997*, pages 192–202, June 1997.

[6] W. Leland, M. Taqqu, W. Willinger, and D. Wilson. On the self-similar nature of ethernet traffic (extended version). *IEEE/ACM Transactions on Networking*, 2(1):1–15, 1994.

[7] O. Rose. Statistical properties of MPEG video traffic and their impact on traffic modeling in ATM systems. volume 20, pages 397–406, 1995.

[8] D. Turaga and T. Chen. Hierarchical modeling of variable bit rate video sources. In *Proceedings of Packet Video 2001*, April 2001.

[9] A. VBavier, B. Montz, and L. Peterson. Predicting mpeg execution times. In *Proceedings of SIGMETRICS 1998*, June 1998.

Title	GOP type	DVD Date	Studio Company	Length (min)	Bit rate (Mbps)	Title	GOP type	DVD Date	Studio Company	Length (min)	Bit rate (Mbps)
Bad Boys	3	6/00	Columbia	119	3.72	Pretty Woman	6	2/00	Touchstone	125	6.96
Barney Super Singing Circus	4	5/00	Lyrick	49	6.45	PrincessBride	3	7/00	MGM	98	5.27
Barney - Barney's House	5	7/99	Lyrick	50	5.91	PulpFiction	5	5/98	Disney	154	6.25
Barney More Songs	3	1/00	Lyrick	49	3.80	Rumble In The Bronx	2	6/97	NewLine	91	5.74
BluesMusicalMovie	5	10/00	Paramount	78	5.02	Rush Hour	4	3/99	NewLine	97	4.98
Book of Pooh	2	7/01	Disney	76	6.41	Rush Hour 2	3	12/01	NewLine	91	4.66
Bug's Life	5	11/99	Disney	95	5.40	Sesame Street 25th Ann	3	7/97	Sony	55	6.18
Caddy Shack	2	4/00	Warner Bros.	99	5.48	Shanghai Noon	2	10/00	Touchstone	110	4.95
Charlie's Angels	3	3/01	Columbia	99	3.93	Shrek	2	11/01	Dreamworks	93	7.05
Chicken Run	3	11/00	Dreamworks	84	4.85	Simpsons Disc 1 Ep. 1	6	9/01	20th Century	25	7.77
City Slickers	4	5/01	MGM	114	4.06	Simpsons Disc 1 Ep. 2	6	9/01	20th Century	25	7.24
Crouching Tiger …	3	5/01	Columbia	120	3.85	Simpsons Disc 1 Ep. 3	6	9/01	20th Century	25	7.25
Die Hard	6	7/01	20th Century	132	6.26	Simpsons Disc 1 Ep. 4	6	9/01	20th Century	25	6.48
Die Hard 2	5	7/01	20th Century	124	6.00	Simpsons Disc 1 Ep. 5	6	9/01	20th Century	25	7.62
Die Hard 3	5	7/01	20th Century	131	6.00	Simpsons Disc 1 Ep. 6	6	9/01	20th Century	25	7.41
Dumbo	2	10/01	Disney	64	5.69	Simpsons Disc 2 Ep. 1	6	9/01	20th Century	25	6.51
Elmo in Grouchland	3	12/99	Columbia	73	4.98	Simpsons Disc 2 Ep. 2	6	9/01	20th Century	25	7.18
Enemy of the State	5	6/99	Touchstone	132	5.85	Simpsons Disc 2 Ep. 3	6	9/01	20th Century	25	6.48
English Patient	5	3/98	Disney	162	6.35	Simpsons Disc 2 Ep. 4	6	9/01	20th Century	25	6.43
Entrapment	5	11/99	20th Century	113	6.47	Simpsons Disc 2 Ep. 5	6	9/01	20th Century	25	6.94
Erin Brockovich	5	8/00	Columbia	132	5.41	Simpsons Disc 2 Ep. 6	6	9/01	20th Century	25	6.41
Family Man	5	7/01	Universal	126	4.27	Simpsons Disc 3 Ep. 1	6	9/01	20th Century	25	7.25
Fantasia	5	11/00	Disney	125	5.32	SlapShot	4	1/99	Universal	123	3.96
Fish Called Wanda	4	3/99	MGM	108	4.69	Sneakers	3	3/98	Universal	125	3.90
Friends Vol1 Episode 1	1	12/00	Warner Bros.	24	5.01	SouthPark Vol1 Ep. 1	1	10/98	Warner Bros.	25	5.48
Friends Vol1 Episode 2	1	12/00	Warner Bros.	49	4.98	SouthPark Vol1 Ep. 2	1	10/98	Warner Bros.	25	5.48
Friends Vol1 Episode 3	1	12/00	Warner Bros.	24	4.98	SouthPark Vol1 Ep. 3	1	10/98	Warner Bros.	25	5.48
Friends Vol1 Episode 4	1	12/00	Warner Bros.	24	5.02	SouthPark Vol1 Ep. 4	1	10/98	Warner Bros.	25	5.48
Friends Vol1 Episode 5	1	12/00	Warner Bros.	24	5.51	SouthPark Vol2 Ep. 1	1	10/98	Warner Bros.	25	5.48
Friends Vol2 Episode 1	1	12/00	Warner Bros.	23	6.12	SouthPark Vol2 Ep. 2	1	10/98	Warner Bros.	25	5.48
Friends Vol2 Episode 2	1	12/00	Warner Bros.	22	6.03	SouthPark Vol2 Ep. 3	1	10/98	Warner Bros.	25	5.48
Friends Vol2 Episode 3	1	12/00	Warner Bros.	49	4.98	SouthPark Vol2 Ep. 4	1	10/98	Warner Bros.	25	5.48
Friends Vol2 Episode 4	1	12/00	Warner Bros.	24	5.02	SouthPark Vol3 Ep. 1	1	10/98	Warner Bros.	25	5.48
Friends Vol2 Episode 5	1	12/00	Warner Bros.	24	5.02	SouthPark Vol3 Ep. 2	1	10/98	Warner Bros.	25	5.48
Gladiator	3	11/00	Dreamworks	155	5.17	SouthPark Vol3 Ep. 3	1	10/98	Warner Bros.	25	5.48
The Godfather	6	10/01	Paramount	175	5.44	SouthPark Vol3 Ep. 4	1	10/98	Warner Bros.	25	5.48
The Godfather II disk 1	6	10/01	Paramount	200	6.96	SouthPark Vol5 Ep. 1	1	10/98	Warner Bros.	25	5.48
The Godfather III	6	10/01	Paramount	170	5.43	SouthPark Vol5 Ep. 2	1	10/98	Warner Bros.	25	5.48
Hunt for Red October	4	12/98	Paramount	135	5.98	SouthPark Vol5 Ep. 3	1	10/98	Warner Bros.	25	5.48
Independence Day	6	6/00	20th Century	153	5.21	SouthPark Vol5 Ep. 4	1	10/98	Warner Bros.	25	5.48
IndigoGirls - Watershed	3	1/98	Sony	75	4.70	South Park - the Movie	6	11/99	Paramount	81	5.07
James Taylor Live at Beacon	3	10/98	Columbia	109	4.30	Space Balls	3	4/00	MGM	96	4.27
Jurrasic Park	6	7/01	Universal	127	4.98	Spy Kids	4	9/01	Dimension	88	6.98
Jurrasic Park - Lost World	6	7/01	Universal	129	4.64	Stargate	2	10/01	Artisan	119	6.31
Lethal Weapon 4	2	12/98	Warner Bros.	127	3.62	Stuart Little	3	4/00	Columbia	85	3.73
Mary Poppins	3	3/98	Disney	139	3.97	The Fugitive	2	8/97	Warner Bros.	131	3.84
Matrix	3	9/99	Warner Bros.	136	4.68	Tigger Movie	6	8/00	Disney	77	6.49
Men in Black	3	9/00	Columbia	98	3.91	Tokyo Babylon	1	12/98		50	5.77
Mission Impossible	4	11/98	Paramount	110	4.06	Tokyo Babylon	1	12/98		50	5.01
Mission Impossible 2	5	11/00	Paramount	123	4.16	ToyStory1	6	10/00	Disney	81	5.38
Monty Python - Holy Grail	3	10/01	Columbia	90	5.24	ToyStory2	6	10/00	Disney	92	5.42
Mulan	6	11/99	Disney	88	4.17	True Lies	5	5/99	20th Century	141	6.28
NBA100GreatPlays	3	4/01	Polygram	72	5.78	You've Got Mail	2	5/99	Warner Bros.	120	4.18
Phish Bittersweet Motel	2	3/01	Image	84	5.49						

Figure 9: Compressed Video Library Listing