Video-on-Demand over Satellite

Kobus ROUX[1]

[1]*Meraka Institute, CSIR, POBox 395, Meiring Naude Street, Pretoria, 0001, South Africa*
Tel: +27 12 841 4460, Fax: + 27 12 841 4720, Email: kroux@csir.co.za

Abstract: In countries where cable infrastructure is lacking, satellite communications is often the preferred alternative to reach the mass market. Interactive video-on-demand requires a large amount of bandwidth per subscriber, making it difficult to scale up to many subscribers in a satellite environment. In this article a method and system for interactive video-on-demand over satellite is proposed and evaluated.

The Internet, mobile phone and other communication networks allow the modern business to receive feedback on its products and services more timely and interactively than was traditionally possible. This has lead to the personalisation of websites, marketing campaigns and profiling of users through personalisation techniques such as collaborative filtering. It is shown that by incorporating personalisation into the video-on-demand business model, it becomes possible to grow the number of users independent of the available bandwidth.

Keywords: Video-on-demand, Personalisation, Collaborative Filtering, Bandwidth, Satellite

1. Introduction

A video-on-demand business serves its users by playing movies when requested by users. A new architecture for video-on-demand systems is proposed where personalisation techniques are integrated into the business processes of the video-on-demand business.

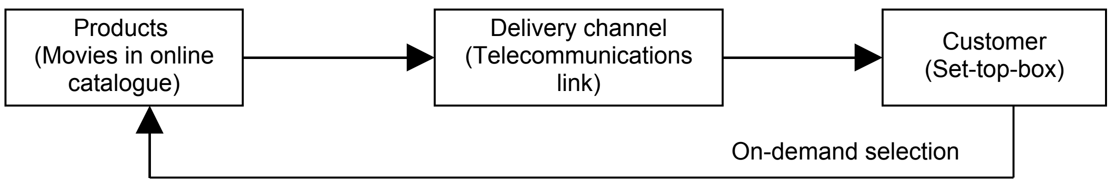

Figure 1: Traditional video-on-demand business model

A diagram of the traditional video-on-demand business model is shown in the above figure. The 'products' are all the movies that are available to the user on the system. The delivery channel is a telecommunications link between the video-on-demand warehouse and the user premises, and could be a satellite, terrestrial wireless or cable link. A user would select items from the catalogue interactively and dynamically (on demand) and expect a movie to start playing soon after it has been purchased.

Any such telecommunications link has a bandwidth limitation, which means that only a limited number of user requests can be serviced simultaneously. A number of compromise approaches to this limitation in video-on-demand systems have been tried, but none has addressed the bandwidth limitation successfully. Two popular approaches in satellite based video-on-demand include:

1. Pseudo video-on-demand, where requested movies start playing only at a certain time; and

2. Time-shifted video-on-demand where multiple time-shifted versions of the same movie is played continuously.

Considering these attempts at solving the video-on-demand bandwidth problem, and considering the use of personalisation, it is proposed that anticipating user requests and downloading the items in advance into a buffer close to the user can significantly reduce the load on the telecommunications link, and create a video-on-demand business that can grow independently of the available bandwidth.

According to current practice, a video-on-demand system requires an amount of bandwidth proportional to the number of users on the video-on-demand system. The growth of a business based on video-on-demand is therefore constrained by the cost and availability of bandwidth. If a recommender system can be used to overcome the bandwidth constraint in a video-on-demand system, then recommender technology can also have application in other businesses.

2. Objective

This objective is to develop a system that can provide interactive video-on-demand, i.e. allowing the user to interactively select movies from a catalogue and the system then play the movie. What is important to the user of a video-on-demand business is the range of titles available (size of catalogue), the responsiveness of the system (how long it takes to start playing a movie) and the locality of the service, e.g. having it available at home.

3. System design

Considering the generic video-on-demand model, one can identify that the key constraint is the available bandwidth on the telecommunications link. (See Figure 1)

Current approaches to addressing this problem were based on removing duplication in playing of movies over the same link. Another different approach to achieve this would be to have a local buffer at the user premise, and then broadcast each movie only once to store it in the local buffers. When the user then requests a movie, it is played from the local buffer and not from the head-end. This approach is shown in the figure below.

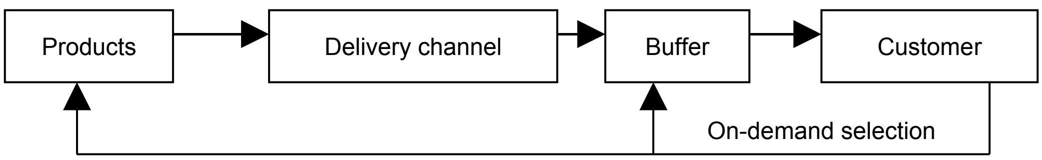

Figure 2: Proposed Video-on-Demand architecture

The problem then shifts to the buffer, as this would have to be quite large to store all the movies in the catalogue. It is therefore proposed that anticipating user requests and storing only the system recommended items in advance into a local buffer can significantly reduce the load on the telecommunications link.

Using a recommender system based on collaborative filtering techniques, the system can predict which movies the user is likely to request. Having this knowledge in advance allows the system to download the appropriate movies to the buffer in the background using available bandwidth during quiet times, and then still providing a true video-on-demand experience for the user when selecting those movies to watch.

4. Recommender system technology

The concepts of collaborative filtering and recommender system technology are well known in the fields of marketing and customer relationship management. Probably the most well known example of a business using collaborative filtering is Amazon.com, which recommends items as you browse and select items from their web site. Personalisation techniques have not been used in business processes, due to the potential risks of failure [1].

This is an example of how recommender technology would work: If we want to give recommendations on movies, we require a list of people and the movies they like as shown in the following table.

Table 2: User-item ratings matrix

	StarWars	Batman	Harry Potter	Matrix	Atlantis	Whispers
Ann	✓	✓				
Howard	✓	✓			✓	
Dan		✓	✓			
Rob		✓	✓	✓		
Mary		✓			✓	✓
Jenny	✓	?	?	?	?	?

If we know that Jenny likes StarWars, what else might she like? Batman will not be a good suggestion as everyone likes it and Jenny might already have seen it. However, Atlantis would be a good suggestion, since Howard who also likes StarWars likes it. Furthermore, we might even suggest Whisper, since Mary who likes it also likes Atlantis, which is liked by Howard. Real world problems are much more complex than the example above. We need software applications to help us reach the same goal.

5. System goal

The purpose of the proposed model is to limit the use of bandwidth and optimise the local buffer.

The system generates a list of Top-N recommended movies for the user, based on 'high' and sometimes 'medium' recommendations. The next table indicates how the 'effectiveness' of the predictions can be evaluated. Konstan et al [2] developed a utility function for movie recommendations and stated, "For watching a movie, the value of finding desirable movies is high to movie fans, but the cost of missing some good ones is low since there are many desirable movies for most movie fans. The cost of false positives is the price of the ticket plus the amount of time before the watcher decides to leave." In the case of the video-on-demand system, the value of finding desirable movies is high; the cost of missing good ones is low; and the cost of false positives is both the cost for the 'rental' as well as the potential cost in bandwidth of subsequent requests for movies not on the Top-N recommended list.

Table 1: Effectiveness of system recommendations

		System Recommendation			
		Low	Medium	High	Not recommended
User rating	Low	Good	Acceptable	Bad	Good
	Medium	Good	Good	Acceptable	Acceptable
	High	Acceptable	Acceptable	Very good	Acceptable
	Not rated	Don't care	Don't care	Acceptable	Don't care

The measure of effectiveness is a combination of the number of very good items as 'effective' and the number of bad items as 'ineffective'. As the goal of the complete system is to optimise the bandwidth usage, the ultimate measure is the probability that a user will not request an item from the local buffer, loaded with items in the users' top-N recommended list. This can be calculated as one minus the fraction of bad items plus 50% of unknown items (movies not rated) in the list to the size of the list (N).

Effectiveness = 1 – [Count (Bad items) + 0.5 x Count (Unknown items)]/N … (1)

This effectiveness measure will vary from 0 to 1. If the user rated all the items on the Top-N list 'low', then the probability is 100% that the user would request something that is not on the Top-N list, i.e. effectiveness equal to 0. If the there are 'very good' ratings, then the chance of the user being satisfied with one of the 'very good' items will subtract from the chance of requesting something from the main store, i.e. effectiveness approaching 1. If the system recommends only 'unknown' items, then the chance is 50/50, i.e. 50% chance that the user will still request items from the main store (effectiveness of 0.5).

For the system to work there must be at least one item on the user's list, which will satisfy the user. Other measures, such as differential pricing of local items versus main store items can motivate the user towards selecting a local item, but that still requires at least one valid item for every user.

6. Sample case: Satellite based video-on-demand

Satellite based video-on-demand is a specific implementation of the generic video-on-demand with particular bandwidth limitations. The maximum available bandwidth for a single satellite carrier can be assumed to be 45 MB/s. An interactive download (on demand) from the head-end would require up to 6 MB/s bandwidth. This means that no more than seven users can request an item from the head-end within the same time period (100 minutes on average for a movie).

If one further assumes that users rent on average 4 movies per month, and are likely to watch them at specific times, one can calculate the probability of multiple users requesting items simultaneously.

For the test data, one can select the worst case to be all the users watching at the same time. Thus, no more than seven users should have none likeable items available on their list. Because the a typical cache can be updated over a five day period, and people are unlikely to watch (interactively request) more than 1 item within the 5 day period, it can be assumed that one likeable item is adequate.

7. Test data

For testing of the proposed business model, data is taken from the MovieLens web-based movie recommender (www.movielens.org). The data were sampled from the data collected over a seven-month period from 19 September 1997 to 22 April 1998. The data consisted of 100,000 movie ratings from 943 users on 1682 items. Each user sampled had rated at least 20 items. The data is freely available from www.grouplens.org and researchers are encouraged to use this data set for comparability in results.

8. Evaluation

The proposed business model is tested using the data collected and the following variables:

Let A = number of users
Let K = standard video bandwidth
Let N = number of items stored in the cache (Top-N list)

The model effectiveness is directly related to the bandwidth required for the video-on-demand operation. The bandwidth required can be calculated as the sum of bandwidth required for each user.

$$B = K\sum_{A} BW(n) \quad \ldots (2)$$

In the worst case, this would be equal to the peak number of simultaneous requests times the video bandwidth (K). As any one user could only request or view one movie at a time, the peak number of simultaneous requests is equal to the number of users (A). For evaluation of the effectiveness of the proposed business model, three cases are selected.

The first case is the ***generic*** video-on-demand model. In this case the bandwidth required for every user is independent of other users, and the peak bandwidth requirement equal to:

$$B_1 = K \cdot A \qquad \ldots (3)$$

In the second case, the ***N most popular movies*** are buffered at the user. The bandwidth required equals zero for users who will select an item from the buffer. The peak bandwidth requirement is therefore equal to:

$B_2 = K \cdot A_2$ where A_2=Number of users with no acceptable items in the buffer … (4)

In the third case, the ***recommended*** movies are buffered at the user, as per the proposed business model. The peak bandwidth will then be equal to K times the number of users for whom none of the recommendations are suitable.

$B_3 = K \cdot A_3$ where A_3=number of users with no acceptable items in the buffer … (5)

9. Results

The following table and chart shows the results obtained for the 100 most active users (A = 100) selected from the test data, and assuming a bandwidth requirement of 1 MB/s per movie request (K = 100).

Table 3: Bandwidth required

	Buffer = 1 movie	Buffer = 2 movies	Buffer = 3 movies	Buffer = 10 movies
Generic system (no caching used)	100 MB/s	100 MB/s	100 MB/s	100 MB/s
Caching the most popular and testing against user rating	52.8 MB/s	31.4 MB/s	26.7 MB/s	11.3 MB/s
Caching the most popular and testing against movies seen	43.3 MB/s	23.4 MB/s	18.9 MB/s	6.4 MB/s
Caching the recommended movies and testing against user rating	31.9 MB/s	19.0 MB/s	13.9 MB/s	3.0 MB/s

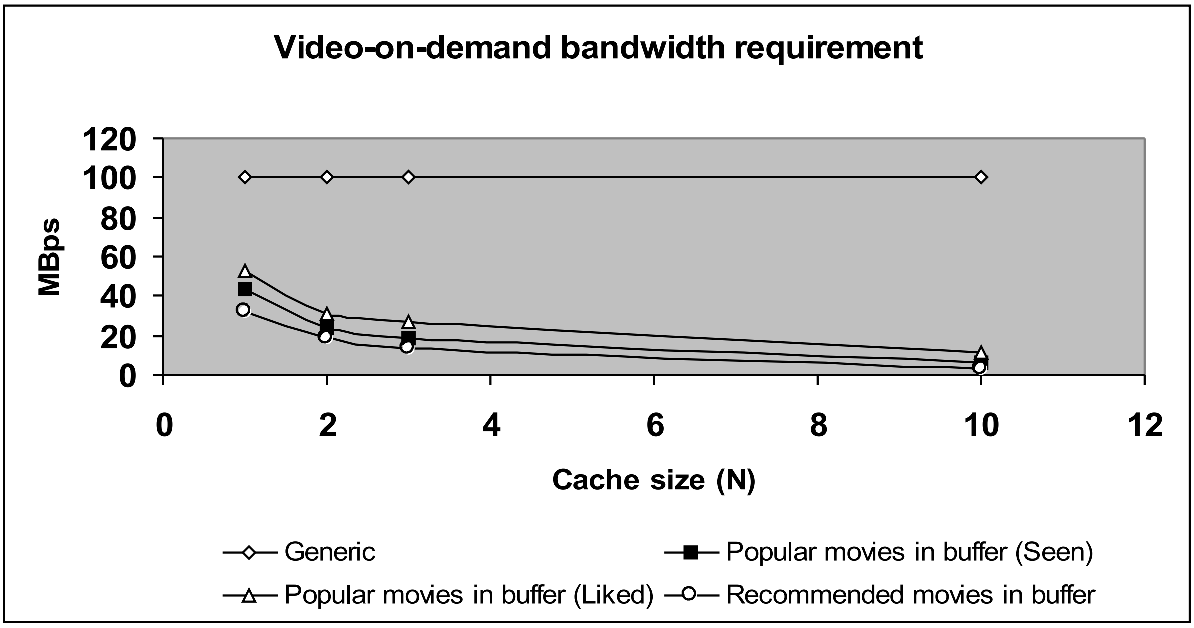

Figure 3: Bandwidth requirement

10. Conclusion

As can be seen from the results above, the proposed model is effective in that it relieves the bandwidth constraint in satellite video-on-demand with up to 97% (3.0 MB/s required with caching, versus 100 MB/s required without caching). This provides anecdotal evidence of the effectiveness of the proposed model.

Video-on-demand may have application in many of the Information Society Technologies (IST) thematic priorities identified by the European Commission, especially in technology enhanced learning.

Personalisation technology is a fairly new technology that has evolved with the large amounts of data and information made available with the Internet. It is an important for researchers in Africa to be aware of this technology and understand its potential value in overcoming the bandwidth constraints.

The utility of personalisation technology in business processes depends on the availability of data and the nature of the business. In this project it has been demonstrated that personalisation has potentially significant value for a video-on-demand business, in providing a method of overcoming the typical bandwidth and scalability problems associated with video-on-demand systems.

While automated collaborative filtering has proven to be accurate enough for entertainment domains (Shardanand & Maes [3]; Konstan et al [2]; Resnick, Iacovou, Sushak, Bergstrom & Rield [4]; Hill, Stead, Rosenstein & Furnas [5]), research has yet to be successful in content domains where higher risk is associated with accepting a filtering recommendation. The main reason for the perceived slower uptake of collaborative filtering is because collaborative filtering systems are stochastic processes that compute predictions based on models that are heuristic approximations of human processes. Second, and probably most important, the systems base their computations on extremely sparse and incomplete data. These two conditions lead to recommendations that are often correct, but also occasionally very wrong (Herlocker, Konstan & Riedl [6]). Herlocker further showed that satisfaction with a recommender system is only partly determined by the accuracy of the algorithm behind it. By explaining and exposing the recommendation process to the user and interacting more with the user, the user satisfaction increased significantly.

11. References

[1] Greeling, D., 2003. “Has collaborative filtering died?”. Discussion on collaborative filtering mailing list (collab@sims.berkeley.edu).

[2] Konstan, J.A., Miller, B., Maltz, D., Herlocker, J.L., Gordon, L., & Riedl, J.T., 1997. Grouplens: Applying collaborative filtering to Usenet news. Communications of the ACM, Vol. 40, No. 3:77-87.

[3] Shardanand, U., & Maes, P., 1995. Social information filtering: Algorithms for automating 'word of mouth'. In Proceedings of the ACM CHI’95 Conference on Human Factors in Computing Systems, pp. 210-217.

[4] Resnick, P., Iacovou, N., Sushak, M., Bergstrom, P., & Riedl, J., 1994. GroupLens: An open architecture for collaborative filtering of Netnews. Paper presented at the ACM

Conference on Computer Supported Collaborative Work (CSCW '94), Research Triangle Park, NC: pp. 175-186.

[5] Hill, W., Stead, L., Rosenstein, M., & Furnas, G., 1995. Recommending and evaluating choices in a virtual community of use. Paper presented at the 1995 Conference on Human Factors in Computer Systems, Denver, Colorado: pp. 194-201.

[6] Herlocker, J.L., Konstan, J.A.., & Riedl, J.T., 2000. Explaining collaborative filtering recommendations. Paper presented at the ACM 2000 Conference on Computer Supported Cooperative Work, pp. 241-250.