

# Investigating the Performance of Video-on-Demand Systems over WLANs Using Generic Association Control\*

Leonardo Bidese de Pinho and Claudio Luis de Amorim

Parallel Computing Laboratory, COPPE Systems Engineering Program,  
Federal University of Rio de Janeiro, R.J, Brazil  
{leopinho, amorim}@cos.ufrj.br

**Abstract.** The problem of scalable video-on-demand (VoD) delivery has been extensively studied over the past years, in the context of wired networks. As VoD-based applications are gaining more and more attention, there is an increasing need of ubiquitous access for them. Wireless networks can be the answer to this need. However, the system behind the application must be scalable to deal with a growing amount of simultaneous accesses, while guaranteeing QoS. Thus the performance of VoD systems must be revisited in this new context, especially aiming scalability. In this paper, we present a preliminary study of a system-level design for VoD systems that operate over 802.11 networks, where commodity Access Points (APs) work collocated in no-overlapping channels which are orchestrated by a central entity that is responsible for controlling the association of clients to the APs. The simulation conducted demonstrates that our system design, using the generic least loaded first heuristic, can use the full aggregate bandwidth as collocated APs allow to. Despite the effective use of bandwidth, it achieved a low blockage rate only for short-length videos at low arrival rates, and performed poorly when video length and arrival rate grow.

## 1 Introduction

Over the last years, multimedia applications and, in particular, those based on the video-on-demand (VoD) technology, have gained increasing attention. From the network perspective, VoD can be seen as a regular application, due to its typical operation: (i) client requests a video from a list of available titles; (ii) a video stream<sup>1</sup> is transmitted with a throughput equal to the average consumption rate of client's video decoder; (iii) while the client watches the video, bandwidth remains allocated to the client over a time equal to the video length.

---

\* Work partially supported by the Brazilian agencies CAPES, CNPq, and FINEP.

<sup>1</sup> Smoothed through a playout buffer at the client, which allows variable bit-rate (VBR) videos to be treated by the system as constant bit-rate (CBR) for allocation of network's bandwidth, also used to hide network jitter.

Hypothetically, in true VoD systems client requests are immediately serviced and video packets are transmitted from the server to the client through a jitter-free, zero latency network. However, in real systems clients will always experience delays between the request and the beginning of the playback, a.k.a. playback latency, that can be broken in three pieces: service latency - the period of time between a request and the actual scheduling of resources to create the stream to feed the client; network latency - that is imposed by the path between the video source and the client; and prefetch latency - period of time to fill the playout buffer at the clients in order to face variations on network latency and video consumption rate. Assuming that both network and prefetch latencies are at their lowest values, a practical true VoD system can be considered the one who achieves zero service latency.

In the context of wired networks, many work have been done to increase the scalability - the capacity of handling simultaneous clients - of VoD systems. In particular, peer-to-peer VoD designs based on stream reuse techniques have shown to be highly scalable for switch based LANs [1], which offer plentiful bandwidth among clients. For WANs and MANs, especially when the network technology employed in the access networks restrict the available bandwidth among the clients, proxy-based techniques seems to be the best choice, as they improve the scalability of VoD applications mainly when proxies are placed at the edges of content distribution networks (CDNs) [2].

In parallel with increasing offer of VoD-based applications, there is a growing need of ubiquitous access to them as more and more people use mobile devices capable of video playback (laptops, PDAs, high-end mobile phones, ...) and want to have access to these applications everywhere. Wireless networks can be the answer to this need. However, the system behind the application must be scalable to deal with this increase of simultaneous accesses, while guaranteeing QoS.

Currently, IEEE 802.11 is the best-seller WLAN technology. Its specification defines both the physical (PHY) and the medium access control (MAC) layers for WLANs. For PHY, the main standards are: a, b, and g. As reported in [3], each of these three physical-layer standards supports a multitude of transmission modes - which specifies the data rate, the modulation scheme, and the error control scheme (e.g.,FEC), if any. In this paper, we focus on the high-speed "a" and "g" 802.11 variants, working in Access Point (AP)<sup>2</sup> infrastructure based mode [4].

In MAC layer, the main access mechanisms are: a mandatory contention-based access protocol - the Distributed Coordination Function (DCF) -, and an optional polling-based protocol - the Point Coordination Function (PCF). The former delivers best effort QoS service level, with no service guarantee in terms of bandwidth. This is true because DCF tries to avoid collisions but does not guarantee they will not occur. On the other hand, PCF can deliver a predictable service performance because the wireless stations are allowed to transmit only when they receive polling messages from the Point Coordinator (PC), which

---

<sup>2</sup> The AP works like a hub, handling all the traffic of a wireless channel (both internal - among wireless stations in the same channel - and external - working as a gateway to the wired network where the AP is connected).

can be included in the AP. Although PCF potentially generates no collision, it is rarely implemented in currently available devices [5]. Given to the problems that were detected in PCF, the IEEE organization has worked on a MAC enhancement for better QoS support, called the 802.11e.

Regarding to bandwidth, while 802.11g offers only three interference-free, co-existent channels over the 2.4 GHz band, the 802.11a supports thirteen channels that operate over the 5 GHz frequency range. As you can see, it is theoretically possible to join channels' capacities, so that we can have sixteen parallel channels in the coverage area of the "collocated APs"<sup>3</sup>. Although each channel has 54 Mbps of maximum link rate, the maximum effective throughput using UDP is near to 30.7 Mbps, leading to a total aggregate bandwidth of 491.2 Mbps [6]. Given to the characteristics we explained, 802.11 can be seen as a potential candidate for access network in a VoD delivery system, especially in environments where support for mobile and portable stations<sup>4</sup> is a must-have.

Even though the aggregate bandwidth suffices, the 802.11 variants lack a efficient mechanism that optimizes the use of the APs through a procedure, often called "Association Control (AC)" in the literature, for choosing the AP where the client device must be associated when it is in the coverage area of multiple APs - collocated or not - using no-overlapping channels.

The basic heuristic for association control, namely Strongest Signal First (SSF), gives priority of choice to APs based on the Receive Signal Strength Indicator (RSSI). As reported in many works, for instance in [7], SSF often leads to a poor load balance among the APs, and does not provide minimal bandwidth guarantees. In particular, when applied to collocated APs that approach tends to lead to a random behavior as all the APs will be at the same distance to the client. Other heuristic, widely adopted by wireless vendors but using no-interoperable, proprietary protocols, is known as Least Loaded First (LLF), where the client is assigned to the AP with the highest available throughput. Although this approach optimizes the use of APs [7], the lack of interoperability among different vendors restricts the appeal of the LLF solution. Moreover, for applications that need a guarantee of minimal bandwidth, LLF must be extended to offer this feature, which some vendors do.

The motivation of this work comes from the fact that to the best of our knowledge we have not found any significant work that evaluated performance of VoD systems that explore the aggregate bandwidth of collocated APs, neither one that uses extensions to the LLF association control heuristic to provide minimal bandwidth guarantee to support VoD deployment over 802.11 WLANs.

Thus, in this paper we provide the following contributions:

1. We model the association control problem in a video-on-demand system delivering video through WLANs with collocated APs;

---

<sup>3</sup> APs that are positioned at the same point in space, with the same coverage area, using no-overlapping channels frequencies.

<sup>4</sup> A portable station is one that is moved from location to location, but that is only used while at a fixed location. Actually, mobile stations access the LAN while in motion. Both stations' types are supported by 802.11 [4].

2. We propose a new system-level design for a VoD system operating over 802.11 WLANs that can use efficiently the aggregate bandwidth of collocated APs, while providing minimal bandwidth guarantees through reservation mechanisms closely tied to the association control procedure;
3. We present a conclusive evaluation regarding to the scalable performance of our VoD design using the generic LLF association control heuristic, slightly modified to guarantee minimal bandwidth, for different combinations of video lengths and arrival rates.

The remainder of this article is organized as follows. Section 2 briefly describes related works. In Section 3 we formulate the problem of association control in a VoD system. After we present in Section 4 our proposal for a VoD system design over WLANs. Section 5 shows a performance analysis of our VoD design based on simulation results. Finally, we draw our conclusions and highlight our ongoing work in Section 6.

## 2 Related Work

To the best of our knowledge, we have not found any significant work directly related to ours, though there are interesting initiatives, focusing on VoD delivery over WLANs and on strategies to control and optimize WLAN throughput, using association control, channels aggregation, and mesh networks.

For instance, the work in [8] presents a VoD system design - namely MobiVoD - for mobile ad hoc clients, using periodic broadcast to increase the scalability of the VoD system. However, MobiVoD does not explore collocated AP's bandwidth, and reports results for a single video, which is not our case. Another work is the WiVision system [9] that supports both live and on-demand delivery of video over WLANs located in the last-mile. Even though they provide practical results of the system for multi-channel 802.11b networks, they assume a best-effort service and measure errors occurred in different situations, which also is not our case since we support minimal QoS guarantee.

Many works tackle the problem of association control of clients focusing on load balance (aiming at homogeneous usage of APs' bandwidth) and fairness (trying to offer the same QoS to all users). [10] indicates the use of schemes that consider the APs' load at association time, optimizing the usage of the network. Moreover, it suggests that clients reassociation mechanisms must be provided to be used when load unbalancing is detected. [11] proposed a dynamic load balance algorithm, namely DLBA, that performs association control aiming to maximize average RSSI and minimize the variance of both RSSI and the amount of clients per AP. The simulated results, considering that all clients have the same traffic profile, showed that DLBA is able to fairly associate clients to the APs. The work in [12] uses information that is implicitly provided by the 802.11h standard regarding to channels' conditions in the association process. The results showed that its scheme outperformed the SSF heuristic, especially when there was external interference. In particular, the work in [7] proposed a max-min

fairness scheme whose performance outperformed the SSF and LLF heuristics, and that by balancing the load on the APs the overall network throughput could be increased. Although that work does not study the behavior of the system by focusing on application needs, it helped us to better understand the problem of association control and inspired us to apply it in the VoD context. Although our work is directly related to the problem of association control, it can be seen as a new approach, focused on application requirements.

In the other hand, [13] presents a significant example of works focusing on channels aggregation that proposes a scheme where stations capable of operating on multiple channels simultaneously can group channels to increase their own throughput, without blocking other stations that support a single channel to share the same network. Besides academic initiatives, channel aggregation is employed in the new IEEE standard 802.11n and in WLAN equipments based on proprietary solutions. Although channel aggregation increases the throughput per channel effectively, it reduces the number of non-overlapping channels, so that the aggregate bandwidth (which is at the core of our work) that can be offered is not significantly affected.

Also, many works have reported significant improvements on the throughput of ad hoc networks organized as meshes, in particular those using either multiple channels [14] and/or interfaces [15]. Currently, our work used infrastructure based WLANs, but it can be extended to tackle mesh networks too.

### 3 Association Control Problem

In this section, we model the association control problem for a VoD system that delivers video through WLANs, where there is a set of clients with mobile devices ( $MD$ ), each of which has a multi-band 802.11 interface that can be configured to either "a" or "g" through software.

Let  $Max_{channels}$  be the maximum number of no-overlapping channels. There is a set of collocated APs ( $AP_{total}$ ), where  $1 \leq AP_{total} \leq Max_{channels}$ .  $MinAP_{throughput}$  and  $MaxAP_{throughput}$  are the minimum and maximum effective AP throughput, respectively. The effective throughput ( $AP_{throughput_i}$ ) is the throughput (in Kbps) of the  $i$ th AP, so that  $MinAP_{throughput} \leq AP_{throughput_i} \leq MaxAP_{throughput}$ . Thus we can derive the Aggregate Bandwidth ( $AB$ ) of the wireless segment of the system, which is equal to  $\sum_{i=1}^{AP_{total}} AP_{throughput_i}$ .

We consider a set of videos ( $V_{total}$ ), where  $MinV_{rate}$  and  $MaxV_{rate}$  are the minimum and the maximum video rate - the average consumption rate (in Kbps) of the video at the decoder -, respectively. Moreover, let the minimum and maximum video length - duration of the video (in sec) - be  $MinV_{length}$  and  $MaxV_{length}$ . Thus the  $i$ th video has a video rate ( $V_{rate_i}$ ) and a length ( $V_{length_i}$ ), so that  $MinV_{rate} \leq V_{rate_i} \leq MaxV_{rate}$  and  $MinV_{length} \leq V_{length_i} \leq MaxV_{length}$ . The system server has an effective throughput ( $S_{throughput}$ ) that can deliver at least  $\frac{S_{throughput}}{MaxV_{rate}}$  and at most  $\frac{S_{throughput}}{MinV_{rate}}$  simultaneous video streams.

We assume that the client request rate (a.k.a. arrival rate) follows a Poisson process with a given  $\lambda$  so that we have in average  $\lambda$  video requests per minute.

Also, the video popularity follows a general Zipf distribution with  $\alpha$  skew. For a given period of simulation (in sec) - Time of Simulation ( $TS$ ) - the total amount of video requests ( $R_{total}$ ) is in average  $\frac{TS * \lambda}{60}$ . Let  $C_{video_i}$  be the id of the selected video for playback by the  $i$ th client. The AP where the  $i$ th client is associated with is  $C_{ap_i}$  and the time it requests  $C_{video_i}$  is expressed as  $T_{request_i}$ , while  $T_{service_i}$  is the time when it starts receiving the stream. From these equations, we derive the Service Latency ( $SL_i$ ) of the  $i$ th client (in sec), which is equal to  $T_{service_i} - T_{request_i}$ . As described before, the system uses a playout buffer at the client, which must be filled up before the playback start. We define this time ( $Prefetch_i$ ) as the prefetching period of the  $i$ th client (in sec). Thus the total amount of waiting time between the request and playback start of the  $i$ th client, the Playback Latency ( $PL_i$ ), is calculated by  $PL_i = SL_i + Prefetch_i$ .

Finally, let  $R_{accepted}$  and  $R_{denied}$  denote the amount of requests that are accepted and denied, respectively. We aim to minimize  $R_{denied}$  to increase the scalability of the system. A new request from the  $i$ th client is accepted only if the necessary bandwidth ( $V_{rateC_{video_i}}$ ) can be allocated in one of the APs. If not, the request is denied. Intuitively, we can say that the availability of bandwidth is a function of the  $\lambda$ ,  $V_{length}$ ,  $V_{rate}$ , and  $AB$ . Thus  $R_{accepted}$  and  $R_{denied}$  will vary according to the combination of these parameters.

## 4 VoD System Design for WLANs

In this section, we present the design we propose for video-on-demand systems operating over 802.11 WLANs. First we overview the hardware and software design, and describe briefly its components. Afterwards, we describe in details the association control procedure.

### 4.1 Overview

From the hardware perspective, the VoD system design (Fig. 1(a)) we proposed can be seen as cluster of collocated 802.11 APs interconnected by a switch with a Gigabit Ethernet port linked to the system server.

The operation of our VoD system involves three software components (Fig. 1(b)), that are responsible for the following specific tasks:

- Video Player (VP) task: Runs in the client device (laptops, PDAs, high-end mobile phones, and so on with 802.11 multi-band interfaces) for requesting and decoding the selected video for playback;
- Video Server (VS) task: Runs in the system server, handling video requests as a real server (being the main storage of video content) or as a proxy.
- Wireless Channel Manager (WCM) task: Can run by default in the system server or in another machine logically connected to the switch with acceptable latency. WCM is responsible for the association control, by choosing the AP where the client must be associated with.

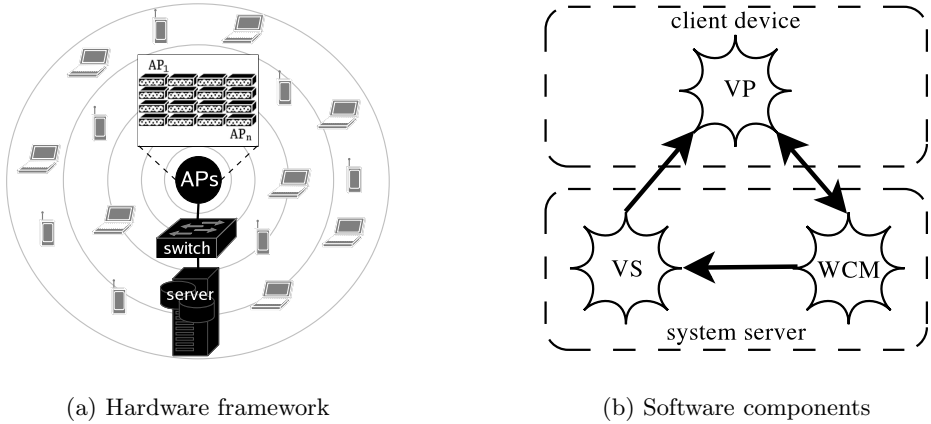


Fig. 1. VoD system design for WLANs

## 4.2 Association Control Procedure

Every time a client selects a video for watching, its video player sends a request token - in the form of  $(CID, VID)$ , where  $CID$  is the client id and  $VID$  the video id - to WCM through any one of the APs. Depending on the availability of bandwidth of the collocated APs, the WCM accepts or denies the video request. When accepted, WCM sends the token  $(CID, APID)$  to  $CID$ , where  $APID$  is the id of the AP that  $CID$  must be associated with, and signals VS to start a stream of the video  $VID$  to  $CID$ . If denied, a  $APID=NULL$  is sent. After receiving the last video packet, the client sends a release token  $(CID, VID=NULL)$  to WCM, freeing the bandwidth.

The WCM implements the Least Loaded First (LLF) [7] association control heuristic, in a modified manner to guarantee minimum bandwidth. To provide this service, WCM maintains a vector -  $FreeBandwidth[n]$  - to keep the available bandwidth of each one of the  $n$  collocated APs, and uses the following algorithm to either accept or deny requests:

1. For  $i$  from 1 to  $AP_{total}$  do
  - Initialize  $FreeBandwidth[i]$  with  $AP_{throughput_i}$
2. For each token  $(CID, VID)$  received do
  - If  $VID=NULL$  then do
    - Decrements  $V_{rateC_{videoCID}}$  from  $FreeBandwidth[C_{apCID}]$
  - Else set  $APID=NULL$  and for  $i$  from 1 to  $AP_{total}$  do
    - If  $V_{rateVID} \leq FreeBandwidth[i]$  then do
      - \* Set  $C_{videoCID}=VID$ ,  $C_{apCID}=i$ , and  $APID=i$
      - \* Increments  $FreeBandwidth[i]$  with  $V_{rateC_{videoCID}}$
    - Send token  $(CID, APID)$  to  $CID$

## 5 Experimental Analysis

In this section, we investigate the performance of our VoD system using the LLF association control heuristic. In the following subsections, we present the evaluation methodology, the assumptions made for the simulation environment, and then report the performance results we obtained.

### 5.1 Evaluation Methodology

To evaluate the proposal we have built a C++ program working as a discrete event simulator. We used two main metrics to quantify the performance of the system. First, the Blockage Rate ( $BR$ ), which is a metric that hints at the scalability of a VoD system.  $BR$  is defined as follows:

$$BR = \frac{R_{denied}}{R_{total}} \quad (1)$$

where  $R_{total}$  is the total amount of video requests the VoD system received and  $R_{denied}$  is the number of requests that could not be serviced by the WCM due to unavailable resources. Thus, the lower is the BR, the higher is the VoD's system scalability.

Second, the Occupation Rate ( $OR$ ), which demonstrates how efficiently the system uses the aggregate bandwidth.  $OR$  is given by:

$$OR = \frac{PB}{AB} \quad (2)$$

where  $PB$  is the peak bandwidth use and  $AB$  is the aggregate bandwidth that the collocated APs offer, as described in the problem formulation section. Thus, the closer is  $OR$  to one, the better is the use of bandwidth resources.

### 5.2 Simulation Environment

We envision a scenery where thousands of potential video clients have a mobile device ( $MD$ ) - and are grouped in an area smaller than two hundred square meters. For instance, it will be the case for sports events in stadiums, airports, among other examples. Thus we have defined the size of the simulation area equal to 200 m<sup>2</sup>, and divided the area in 5 m<sup>2</sup> regions.  $MD$ s are randomly placed in these regions. Each  $MD$  has its own id, so that the first client that requests a video has id=1, the second has id=2, and so on. As said before, we assume that the arrival of client requests follows a Poisson process with  $\lambda$  and the choice of videos follows a Zipf distribution with  $\alpha = 0.7$  [16].

Moreover, our simulation assumes the following assumptions:

- Contention-free Dedicated Network: the network's bandwidth is dedicated only to the traffic that the VoD system generates. Also, as we have a centralized main source of traffic - the video server - we assume that in practice the contention at the APs will be insignificant given to the small size and



amount of control packets that are exchanged among the units. Thus, the system can use DCF without standard modifications, while guaranteeing QoS to the application;

- Transmission Range and Standards: we consider the coverage area of the APs equal to the simulation area and the use of 802.11a/g-compliant APs;
- Channel Interference: we adopted only non-overlapping channels, and did not investigate interference from other types of signal sources. Regarding to the use of multiple APs near to each other, we assumed that the interference among them will be in an acceptable threshold provided that each AP is at least 60 cm far from any other AP, as shown in [14];
- Transmission errors: although error resilience is an important issue in wireless video delivery, our current proposal did not tackle such a problem. However, note that VoD allows trivial error treatment at both the network and the application level;
- Equalized Signal-Strength: all APs are located in the same point in the simulated scenario, using the same transmission power. Thus, the signal level is the same for all APs, as perceived by the *MD*. As there is no strongest-signal AP, the SSF association control heuristic is useless in this context.

### 5.3 Performance Evaluation

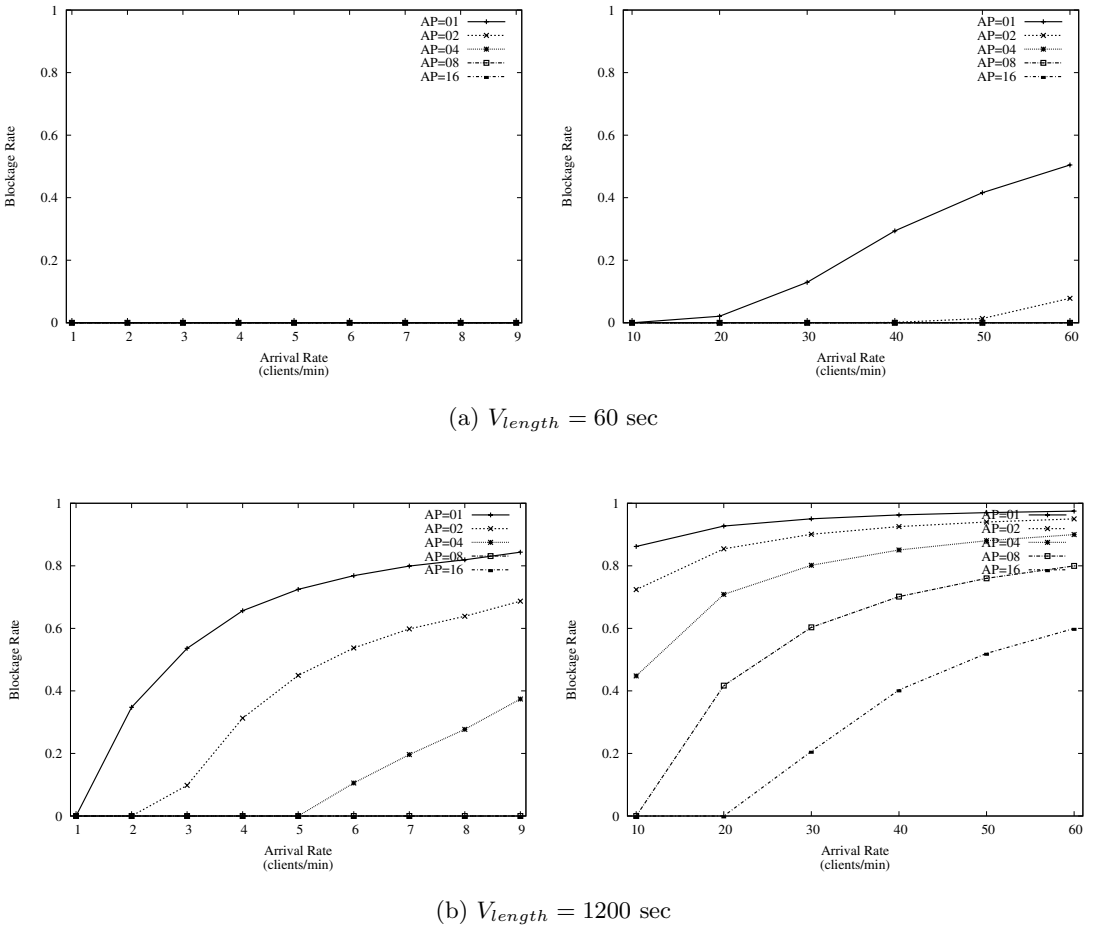
In this subsection, we describe the simulation results. Given the restricted space, we will concentrate our analysis on videos of same length and rate, ranging from one to twenty minutes, which we assume will be most popular for environments with mobile and portable devices, especially because wireless devices have battery constraints. Also, we restricted the analysis on APs with the same throughput. Table 1 summarizes the parameters of the simulation.

**Table 1.** Simulation Parameters

| Parameter                                  | Values         |
|--|----------------|
| Simulation Time - $TS$ - (sec)             | 3600           |
| Collocated APs - $AP_{total}$ - (un)       | 1, 2, 4, 8, 16 |
| AP Throughput - $AP_{throughput}$ - (Kbps) | 30720          |
| Video Length - $V_{length}$ - (sec)        | 60, ..., 1200  |
| Video Rate - $V_{rate}$ - (Kbps)           | 1024           |
| Videos - $V_{total}$ - (un)                | 100            |
| Arrival Rate - $\lambda$ - (requests/min)  | 0.1, ..., 60   |

Note also that, even though we are using multiple videos, the simulated behavior would be the same for a single video because they have the same rate and length. In a future work we will evaluate the results for multiple videos with different combinations of video lengths and rates.

Fig. 2 shows in details the impact of the amount of APs and arrival rate on the blockage rate for two video lengths (60 and 1200 seconds). To better illustrate the results, we divided the graphics by ranges of arrival rates.



**Fig. 2.** Impact of different combinations of amount of collocated APs ( $AP_{total}$ ) and arrival rate ( $\lambda$ ) on the blockage rate ( $BR$ ) for the shortest and longest video length

Another important measurement, derived from  $BR$ , is the minimal arrival rate ( $\lambda$ ), hereafter expressed as  $MAR$ , where the combination of video length ( $V_{length}$ ) and amount of APs ( $AP_{total}$ ) forces the WCM to deny requests for lack of free bandwidth. Table 2 presents the values we have collected for  $MAR$ .

As you can see in the curves of Fig. 2 and in the values of Table 2, a single AP is enough to deliver all the video-lengths analyzed for arrival rates lower than

**Table 2.** *MAR* results: the minimal arrival rate where  $R_{denied} > 0$ 

| $AP_{total}$ (un) | $Video_{length}$ (sec) |     |     |     |     |     |     |     |     |      |      |
|-------------------|------------------------|-----|-----|-----|-----|-----|-----|-----|-----|------|------|
|                   | 60                     | 120 | 240 | 360 | 480 | 600 | 720 | 840 | 960 | 1080 | 1200 |
| 1                 | 20                     | 9   | 5   | 3   | 3   | 2   | 2   | 2   | 2   | 2    | 2    |
| 2                 | 40                     | 20  | 10  | 8   | 6   | 5   | 4   | 4   | 3   | 3    | 3    |
| 4                 | -                      | 50  | 30  | 20  | 20  | 10  | 9   | 8   | 7   | 6    | 6    |
| 8                 | -                      | -   | 60  | 40  | 30  | 30  | 20  | 20  | 20  | 20   | 20   |
| 16                | -                      | -   | -   | -   | 60  | 50  | 40  | 40  | 30  | 30   | 30   |

two clients per minute. As the arrival rate grows, the number of collocated APs must be increased to avoid request denials. For the arrival rates we measured and short-length videos, less equal to 360 seconds, that suffices. However, for longer videos the system had to deny requests. For the highest arrival rate ( $\lambda = 60$ ), the system performed poorly ( $BR > 0.2$ ) for video length longer than 600 seconds.

We have also analyzed the Occupation Rate (*OR*) when  $R_{denied} > 0$ . In this case, the *OR* measured was always one, which proves the effective use of the aggregate bandwidth by the system.

## 6 Conclusions and Ongoing Work

In this paper, we investigated performance of video-on-demand (VoD) systems running over WLANs with collocated access points (APs), using an extended version of a generic association control heuristic to choose the APs the clients use, while providing the minimal QoS that is required by the application.

To achieve this goal, we have modeled the association control problem and proposed a new system-level design for VoD system operating over 802.11 that can use efficiently the aggregate bandwidth of collocated APs, while providing minimal bandwidth guarantees through reservation mechanisms closely tied to the association control procedure.

We also have presented a conclusive evaluation regarding to the scalable performance of the system using the generic association control heuristic LLF, slightly modified to guarantee minimal bandwidth, for different combinations of video lengths and arrival rates. The simulation we conducted demonstrated that our VoD design, using LLF, could use the full aggregate bandwidth that collocated APs allowed. Moreover, the system achieved a low blockage rate for short-length video and low arrival rates, but performed poorly as the video-length and arrival rate grows. However, based on these results, we speculate that a heuristic more application-friendly will significantly reduce such a problem.

Thus, our next step we will be concentrated on investigating new VoD-friendly association control heuristics. Other interesting topic we are aiming at is to study the performance of the system in heterogeneous 802.11 networks, where client devices can have either multi-band or single-band interfaces. In parallel, we have engaged in the TRAVIS-QoS project (<http://www.lcp.coppe.ufrj.br>) where

we aim to build a wireless test-bed and a proof-of-concept prototype that will be used to corroborate or not the results we obtained so far through simulations.

## References

1. Pinho, L.B., Amorim, C.L.: Assessing the efficiency of stream reuse techniques in P2P video-on-demand systems. *Journal of Network Computing Applications (JNCA)* (2004) (article in press).
2. Ishikawa, E., Amorim, C.L.: Collapsed Cooperative Video Cache for Content Distribution Networks. In: *Proceedings of the Brazilian Symposium on Computer Networks (SBRC)*, Natal, RN, Brazil (2003) 249–264
3. Manshaei, M., Turletti, T., Krunz, M.: A media-oriented transmission mode selection in 802.11 wireless LANs. In: *Proceedings of the IEEE Wireless Communications and Networking Conference (WCNC)*, Atlanta, GA (2004)
4. ANSI/IEEE: Wireless LAN medium access control (MAC) and physical layer (PHY) specification, 1999 edition [iso/iec 8802-11: 1999] (1999)
5. Pavon, J., Choi, S.: Link Adaptation Strategy for IEEE 802.11 WLAN via Received Signal Strength Measurement. In: *Proceedings of the IEEE International Conference on Communications (ICC)*, Anchorage, AK (2003)
6. McFarland, B., Wong, M.: The family dynamics of 802.11. *ACM Queue* **1** (2003)
7. Bejerano, Y., Han, S.J., Li, L.E.: Fairness and Load Balancing in Wireless LANs Using Association Control. In: *Proceedings of the International Conference on Mobile Computing and Networking (MobiCom)*, Philadelphia, PA (2004) 315–329
8. Tran, D.A., Le, M., Hua, K.A.: MobiVoD: A Video-on-Demand System Design for Mobile Ad Hoc Networks. In: *Proceedings of the IEEE International Conference on Mobile Data Management (MDM)*, Berkeley, CA (2004) 212–223
9. De, P., Sharma, S., Shuvalov, A., Chiueh, T.: WiVision: A Wireless Video System for Real-Time Distribution and On-Demand Playback. In: *Proceedings of the IEEE Consumer Communications and Networking Conference (CCNC)*, Las Vegas, NV (2004)
10. Raniwala, A., Chiueh, T.: Coverage and Capacity Issues in Enterprise Wireless LAN Deployment. Technical Report TR-166, Stony Brook University, Computer Science Department, Experimental Computer Systems Lab (2004)
11. Sheu, S.T., Wu, C.C.: Dynamic Load Balance Algorithm (DLBA) for IEEE 802.11 Wireless LAN. *Tamkang Journal of Science and Engineering* **2** (1999) 45–52
12. Korakis, T., Ercetin, O., Krishnamurthy, S., Tassiulas, L., Tripathi, S.: Link quality based association mechanism in IEEE 802.11h compliant wireless lans. In: *Proceedings of the 1st workshop on Resource Allocation in Wireless Networks (RAWNET)* (in conjunction with IEEE WiOpt 2005), Riva Del Garda, Italy (2005)
13. Mangold, S., Habetha, J.: IEEE 802.11 contention-based medium access for multiple channels. In: *Proceedings of the 11th European Wireless Conference*. Volume 2., Nicosia, Cyprus (2005) 522–526
14. Raniwala, A., Chiueh, T.: Architecture and Algorithms for an IEEE 802.11-based Multi-channel Wireless Mesh Network. In: *Proceedings of the INFOCOM*, Miami, FL (2005)
15. Zhu, J., Roy, S.: 802.11 Mesh Networks with Two Radio Access Points. In: *Proceedings of the IEEE International Conference on Communications (ICC)*, Seoul, Korea (2005)
16. Dan, A., Sitaram, D., Shahabuddin, P.: Dynamic Batching Policies for an On-Demand Video Server. *Multimedia Systems* **4** (1996) 112–121