Comparing Network Design for the Provision of Distributed Home Theatre Services

Nelson.L.S. Fonseca, Cristiane M.R. Franco
F Schaffa

Relatório Técnico 97-18

Outubro de 1997

Comparing Network Designs for the Provision of Distributed Home Theatre Services

Nelson L. S. Fonseca and Cristiane M. R. Franco
State University of Campinas
Institute of Computing
P.O.Box 6176
13083-970 Campinas SP
Brazil
Phone/fax: +55+19+2530123
e-mail: nfonseca@dcc.unicamp.br
and
Frank Schaffa
IBM T. J. Watson Research Center
Yorktown Heights, NY 10598
U.S.A

Technical Subject Area: Multimedia Communications

Abstract

Video services is both a major business driver and a bandwidth consumer for the future broadband integrated network. Understanding different video services requirements is of paramount importance for network design. In this paper, we study Distributed Home Theatre (DHT), a video service which allows distributed users to debate a film. We compare different network designs for the provision of DHT services, and study the interplay between bandwidth reduction and program replication. We evaluate the impact of user distribution, and number of users per session on this interplay. Moreover, we analyze networks with both DHT and video on demand services

I) Introduction

The advent of traffic integration opened avenues for countless multimedia applications. Among the most promising applications are video-conference and video-on-demand. Video-conferences extends current voice-conference by allowing participants to see their body language and to use visual information. Users of a video-on-demand system can select and watch films from video archives. In this paper, we consider Distributed Home Theatre (DHT), a hybrid application in which a film is simultaneously played for the participants of a video-conference. The basic idea is to allow a group of distributed users to discuss a film. In a DHT session, any participant can initiate a debate about a specific scene by performing VCR operations on the video. It is expected that distributed home theatre will have a great impact on distance-learning as well as on professional conferences [1] [2].

Video applications will be the major bandwidth consumers of the future broadband integrated network. Take, for example, the deployment of video-on-demand service in the continental United State where there are approximately 77,000,000 viewing households during prime-time [3]. Using data rates of 6 Mb/s for MPEG-II NTSC, 10 Mb/s for JPEG NTSC, and 20 Mb/s for HDTV transmission, the total bandwidth requirements would be 462 Tb/s, 770 Tb/s and 1.54 Pb/s, respectively [4]. These requirements are far in excess of the current network infrastructure. Even with the deployment of high bandwidth switches and links, the huge bandwidth consume calls for proper network engineering.

In [5], we analyzed the interplay between bandwidth and replication of stored programs in network design for the provision of DHT services. We introduced a cost function which took into account both server and bandwidth costs, and we evaluated different network scenarios by considering server replication and caching. Moreover, we investigated networks with both Distributed Home Theatre and Video-on-Demand (VoD) services [6]. It was clear that server/cache replication techniques are not enough to reduce the hugh bandwidth requirements of DHT services, and networks with bandwidth reduction characteristics should be considered for the deployment of DHT services in large scale.

In this paper, we investigate the design of networks with stream sharing and with distributed video server for the provision of DHT services. We consider server/cache replication techniques and evaluate the impact of parameters such as the number of users per DHT session, and user geographical distribution on the network cost. We extended our analysis to include networks with both DHT and VoD services. Furthermore, we compare the mentioned networks in providing distributed video services.

This paper is organized as follows: section II introduces distributed home theatre and different network designs. Section III describes design techniques for bandwidth reduction. Section IV introduces a methodology to compute network cost. Sections V and VI analyse networks with stream sharing and networks with distributed server. Finally, conclusions are drawn in section VII.

II) Distributed Home Theatre

The basic idea of Distributed Home Theatre (DHT) is to allow a group of distributed users to discuss a film. DHT can be seen as a video/voice conference in which users are watching a film. Any participant of a DHT session can issue VCR operations and initiate a debate about a specific scene. Distributed Home Theatre is a promising tool for distance learning as well as for professional conferences.

We assume a hierarchical distribution network composed by a national ATM backbone, metropolitan ATM networks and local loops connected to the metropolitan network via head end ATM switches (Figure 1). In the current CATV infrastructure the number of users connected to a single tree must be less than a thousand in order to guarantee minimum Quality of Service [7]. This trend seems to continue in the future of ATM switched network in order to keep head end switches to a reasonable size. Video servers may be attached at any level of the distribution tree.

The implementation of DHT services depends on characteristics such as switching and server architecture. In this paper, we investigate networks with different characteristics in the provision of DHT services. In the simplest case, there is a video stream for each individual user and all streams of a DHT session are distributed by the same server (DHT *without* stream sharing) [5]. In networks *with* stream sharing, a video stream is shared among users of a DHT session up to the deepest common node of a session distribution subtree (a subtree connecting all participants of a DHT session), and all the streams of a DHT session are furnished by the same server. In a third approach, a video stream is provided by the closest server to the user. Only control messages are exchanged by servers involved in a DHT session (DHT with distributed server). These three approaches represent different compromise between bandwidth consuption and bandwidth demand on individual server.

We assume that the distribution network is a balanced d-ary tree (Figure 2). Each node of the tree represents a switch and the leaves represent head ends switches which connect users to the network. Each switch may have a server (or a program cache) attached to it. The number of head end H is determined by the number of users with a maximum number of users per head end. The number of hierarchy levels is given by:

$$L(d) = \log_d(H)$$

and the number of switches $W(d)$ in such network is given by:

$$W(d) = \sum_{n=1}^{L(d)} d^n$$

Although in the numerical examples we use a binary balanced tree with 1024 head ends and 1000 users per head end, our methodology can also be applied to non-balanced d-ary trees.

III) Designing Techniques

The trade-off between bandwidth and program replication guides the design of networks with video services. If bandwidth costs were negligible, we would have a central server (or servers) providing service to all network users. Conversely, if storage cost were close to zero, we would have a video archive at every user set top box. Obviously, none of these approaches are realistic. The current network resources are far behind the huge bandwidth demand generated by the central server solution, and the video server architectures are much more complex and expensive than the current personal computers. Consequently, a solution to ameliorate the bandwidth requirements consists in replicating a certain number of servers in some nodes of the distribution tree. In a real network design, the number and locations of servers are determined by regional demand, network topology and current cost of technology.

In our attempt to understand the trade off between bandwidth and replication of storage program, we analyze a server replication strategy. We initially consider the cost of a single server located at the root of the distribution tree (level zero). In the second step, we place servers only at (all) nodes of level 1. We proceed by considering at step $l+1$, networks with servers located only at the l^{th} level of the distribution tree.

We also investigate a program caching strategy [8]. Instead of replicating the whole server, we replicate just the most popular programs. In this way, we try to reduce the overall cost by reducing individual server costs. However, we need to provide a full server at the root of the tree to handle requests to non-popular programs (cache misses). To evaluate the impact of adopting program caching, we use the same rationality used when studying server replication: at step l+1, we consider caches only at the l^{th} level (in addition to the full server at the root).

IV) Network Costs

In this section, we show a framework to compute network costs which can be used for engineering real networks. In our analyses we do not take into account voice and control signals given that video is the dominant component of the bandwidth consume.

The bandwidth cost of a DHT session is determined by the number of allocated links to it. Undoubtfully, user dispersion is a key parameter which impact bandwidth costs. However, user behavior might only be fully understood when service is deployed. Therefore, we evaluate our results considering three different distributions. We use: i) a uniform distribution, ii) a normal distribution, and iii) a normal distribution per head end to represent loosely, moderately and highly concentrated patterns of users per DHT session, respectively. The mean of the normal distribution is equal to the media of the number of head ends and each normal distribution per head end is centered in the associated head end. For instance, when considering a normal distribution per head end with mean equals to 525 and with standard deviation equals to 5, 68.26% of the participants of a DHT session are statistically located between head end 520 and head end 530.

The total bandwidth cost (C_b) is the sum of the cost of each allocated link, and the cost of an allocated link is proportional to the number of programs delivered through it. Thus,

$$C_b = \gamma_b \sum_{\Lambda} b(\lambda_i)$$

where:

Λ - is the set of links,

$b(\lambda_i)$ - is the bandwidth on link λ_i,

γ_b - is a normalization constant.

C_b is normalized by the factor γ_b which makes it easier to explore different cost scenarios.

The cost of a server depends on the number and on the access rate of the programs stored in it. The access rate of a program is related to the bandwidth needed to support the incoming requests. For example, a server with the top-ten most popular (higher access rate) movies may need to support more bandwidth (and consequently, costs more) than a large achieve of unpopular programs.

In order to capture the effect of both factors in the server cost, we use Zipf's law [9]. It was shown that Zipf's law accurately models the popularity of rented movies in United State by using data published in specialized magazines such as Billboard Magazine and Video Store Magazine. To derive our results we assume that the popularity of programs in DHT services will be the same as their popularity in rental stores. Zipf's law says that the probability of choosing program i among N_p stored programs is given by:

$$z(i) = C / i$$

$$C = 1 / \sum_{i=1}^{N_p} (1/i)$$

We assume that the server at the root of the distribution tree costs one unit. Thus, the storage cost at the root is:

$$C_{root} = \sum_{i=1}^{N_p} z(i)$$

where $z(i)$ is the probability of choosing program i where programs are ordered according to their decreasing popularity.

The total server cost is given by the summation of the cost of each individual server:

$$C_s = \sum_{\Delta} C_i$$

where Δ is the network set of servers.

Finally, the total cost to provide DHT services is given by the summation of the bandwidth cost with the server cost.

$$C_t = C_b + C_s$$

In this paper, results were obtained via simulation. We used the replication methods for generating confidence intervals with 95% of confidence level. The sample size used to compute each point of the curves was such that confidence intervals width were less than 5% of the mean. Given that results with a high degree of confidence were generated, we only show the mean value for the sake of visual interpretation.

To compare trends we normalize both bandwidth and server/cache costs by the highest value in each curve. We also display results by normalized depth (level 0 corresponds to the root level and level 1 to the head end level). In our numerical examples, we display the impact of the number of users per session and user distribution on the network cost.

V) Networks with Stream Sharing

In networks *with* stream sharing, a video stream is shared by the participants of a DHT session up to the deepest common node of a session distribution subtree. In this section, we investigate the design of networks *with* stream sharing and compare it with the design of networks *without* stream sharing.

V.1) Sever Replication

In networks *with* stream sharing, the bandwidth cost of a DHT session is the same for every server located along the session distribution subtree. Therefore, when considering the bandwidth cost for server replication at the l^{th} level, we may choose any server which belongs to the session distribution subtree at that level.

The total bandwidth cost considering replication at l^{th} level is given by:

$$C_b(l) = \gamma_b \sum_{\beta} c_i(l)$$

where:

$c_i(l)$ - is the cost of the i^{th} DHT session when considering server replication at level l

β - is the set of all DHT sessions.

γ_b - is a normalization constant

In Figure 3, we show the normalized bandwidth cost as a function of the normalized depth for different user distributions. To understand this figure, we observe that the bandwidth cost of a DHT session is the same irrespective of the control server location along the session subtree, i.e., if the server which furnishes the video streams of a DHT session is located at a node of the session subtree, relocating it to another node of the subtree does not affect the bandwidth cost. For loosely and moderately concentrated users, the distribution subtree of a DHT session usually contains nodes close to the root of the distribution tree. Therefore, the bandwidth cost (per session and total) is almost the same for all levels of the distribution tree. While for loosely and moderately concentrated users, the placement level does not significantly impact the bandwidth cost, for highly concentrated users the optimum level is around two levels above the head-end level (normalized depth of 0.7). Although the number of users per session does not have a significant impact on the server optimum location, as we increase the number of users per session we slightly decrease the bandwidth cost given that we increase the degree of video stream sharing

We can easily understand the impact of introducing stream sharing in a network by comparing these results with findings for networks *without* stream sharing (extensive results for networks *without* stream sharing can be found in [5]). In networks *without* stream sharing, as the distribution of participants becomes more concentrated, the optimum placement level moves from the root level towards the head end level, and the number of participants may influence the optimum placement level value for moderately concentrated distributions.

As the number of a DHT session participants increases, the total bandwidth cost decreases since we may increase the degree of stream sharing. The bandwidth reduction due to the increase of the number of participants per session can be as high as 25% for highly concentrated users.

The total server cost for replication at l^{th} level is given by:

$$C_s(l) = \gamma_s \sum_{j=1}^{d^l} \sum_{i=1}^{N_p} z(i/x_j)$$

where:

x_j - is the j^{th} server demand, i.e., the ratio between the number of DHT session served by control server j and the total number of DHT session

$\gamma_s(l)$ - is a normalization constant for level l

The total normalized cost of server replication at level l is given by $C_b(l) + C_s(l)$. Figure 4 illustrates this computation.

From Figure 5, we notice that for loosely and moderately concentrated users the total normalized cost is almost flat until the normalized depth of 0.3. After this point, it increases sharply. As observed before, for these type of distributions, the bandwidth cost is almost the same irrespective of the server placement level. By replicating servers, we reduce the bandwidth demand on each server (and consequently decreases individual server cost). For levels close to the root of the tree, the cost reduction due to less stringent bandwidth demand on each server tends to be compensated by a higher number of servers. For levels deeper than 0.3, the growth of the number of servers tends to dominate the network cost. For highly concentrated users, the bandwidth cost decreases faster than the server cost increases up to the normalized depth of 0.4 (the optimum placement level). After this point, the server cost dominates the network cost. Additionally, we observe that the number of users per DHT session does not impact the placement level.

For higher number of users per session, the network cost decreases due to bandwidth savings. This trend is more striking for highly concentrated users, and the total cost savings can be as high as 25%. When comparing the total network cost of networks *with* stream sharing with the cost of networks *without* stream sharing for the same set of DHT sessions [5], we notice that savings due to a higher number of users per session is in the range of 40% to 50%. This difference corresponds to the most expensive level for server replication in networks *without* stream sharing. The minimum difference occurs at the optimum placement level of networks *without* stream sharing and it is at most of 27%.

For loosely and moderately concentrated users the total cost curve has a similar shape to the curves of networks *without* stream sharing curves [5]. However, in networks *without* stream sharing bandwidth is the major dominant factor of the total cost. For networks *without* stream sharing and highly concentrated users the optimum placement level (0.7) is closer to the head end level than it is for networks *with* stream sharing.

V.2) Caching

For each DHT session, we assign a control cache at level l, i. e., a cache which will handle all the requests for that session. The bandwidth cost for cache replication at the level l is given by:

$$C_b = \gamma_b\left(\alpha\sum_{\beta} c_b(l) + ((1-\alpha)\times\mu\times h)\right)$$

where:

β - is the set of DHT sessions;

γ_b - is a normalization constant;

$c_b\,(l)$ - is the bandwidth cost when considering cache replication at the l^{th} level only;

μ - total number of participants;

h - tree height;

$\alpha = \sum_{i=1}^{M} z(i)$ - is the cache hit probability

and the cache cost for cache replication at level l is given by:

$$C_c(l) = \gamma_c \sum_{j=1}^{d^l} z(1/x_j) + C_{root}$$

where:

$\gamma_c(l)$ - is a normalization constant;

x_j - is the j^{th} cache demand, i.e., the ratio between the number of DHT sessions served by the j^{th} cache and the total number of sessions

In Figure 6, we plot the total network cost as a function of the normalized depth for different cache size. It is evident that caching gives no savings. The network cost increases as we increase the number of cache (the bandwidth cost is constant). For highly concentrated users, caching is also not worth adopting. The savings due to smaller repositories do not compensate the penalty due to cache misses. We note that the optimum placement level (0.4) is the same for server replication.

From [5], we know that for networks *without* stream sharing, caching is also not an attractive strategy when compared to the cost of server replication. In this networks, for moderately and loosely concentrated users the total network cost is almost constant up to level 0.3. After this point, it sharply increases. For highly concentrated users, the optimum placement level is around

0.7. These differences can be understood by the huge bandwidth requirements of a DHT session in networks *without* stream sharing.

By comparing costs, we notice that the introduction of stream sharing can bring savings of up to 50% on the network cost. Savings are higher for highly concentrated users (up to 50%) than it is for loosely/moderately concentrated users (up to 40%).

V.3) Networks with both DHT and VoD services

Video will be the major bandwidth consumer in the future broadband network. It is essential that we take into account the requirements of different video applications in a real network design. One of the most promising applications is video-on-demand. In a video-on-demand system, individual users can select movies to watch from a video server. Providing DHT services costs more than providing VoD services. This happens because users in a DHT system connect themselves to the server which minimizes the bandwidth cost of a DHT session whereas users in a VoD system connect themselves to the closest server. In other words, for the same number of users the total bandwidth cost of VoD sessions is always a lower bound for the bandwidth cost of a DHT session.

In Figure 7, we plot the total network cost for server replication as a function of the normalized depth for different percentage of DHT users. Both network cost and server cost of each curve are normalized by the respective costs of networks with 50% of DHT users. We notice that the optimum placement level is 0.7 irrespective of user distribution. This behavior is highly influenced by the cost of providing VoD services. We observe that for placement levels close to the root of the tree, networks with low percentage of DHT users cost more than networks with high percentage of DHT users (50%). At a certain level, this trend is reversed and the cost of networks with high percentage of DHT users exceeds the cost of networks with low percentage of DHT users. This is because as we move towards the head end level the bandwidth cost of VoD service continuously decreases and at a certain point the bandwidth cost of DHT services begins to dominate the total cost. Hence, after this point networks with higher percentage of DHT users cost more than networks with lower percentage of DHT users. The trend break point depends on the user distribution. For loosely and moderately concentrated distributions the breakpoint are at normalized depth of 0.1 and 0.3, respectively. For this type of distribution, the DHT bandwidth cost is almost constant irrespective of the network level, and it also dominates the total cost at levels close to the root of the tree. For highly concentrated users, the bandwidth cost of both DHT and VoD services decreases as we move towards the head end level giving a break point at a normalized depth of 0.7

When assessing the impact of introducing stream sharing, we notice that, obviously, savings increases as we increase the percentage of DHT users (we do not consider stream sharing among users of VoD services - piggybacking). For networks with 50% of DHT users savings can be as high as 45%. For networks with 10% of DHT users savings are around 5%.

VI) Networks with Distributed Server

In networks with distributed servers, the closest server to a user furnishes video streams to him/her. Only control messages are exchanged among servers supporting a DHT session. Networks with distributed servers is a clear compromise between reducing bandwidth and increasing the bandwidth demand on video servers (note that we are not considering in this paper distributed server combined with stream sharing). Moreover, this solution requires more complex architecture in order to cope with synchronization issues.

The analysis of networks with distributed servers is identical to the analysis of networks with VoD [4]. Therefore, user distribution does influence the network cost. In this section, we show the main results for networks with distributed servers and we briefly compare it to the findings of networks without stream sharing.

VI.1) Server Replication

A server located at a node of the l^{th} level furnishes video streams to all users connected to head ends which are descedants of that node. Hence, the total bandwidth cost is given by:

$$C_b(l) = (L(d) - l + 1) \times U$$

where:

U - is the number of users (subscribers);

$L(d)$ - is the height of a d-ary tree;

The total server cost for replication at level l is given by:

$$C_s(l) = \gamma_s \left(d^l \times \sum_{i=1}^{N_p} z(i/x_j) \right)$$

where:

$x_j = U / d^l$ - is the j^{th} server demand;

γ_s - is a normalization constant.

Figure 8 illustrates the network cost as a function of the normalized depth. The bandwidth cost continuously decreases as we move towards the head end level. It decreases faster than the server cost increases. After the normalized depth of 0.7 (optimum level) this trend is reversed. Placing the server at deeper levels increases the network cost. Networks *without* stream sharing with highly concentrated users has a similar network cost to networks with distributed servers. The difference between them is that for networks *without* stream sharing the optimum placement level is at a normalized depth of 0.6. The optimum level is basically dictated by the bandwidth cost (the optimum level for bandwidth is also level 0.6). For loosely and moderately distribution the shape of the network cost curve is quite different as explained in section V.1. When compared to the cost of networks without stream sharing, networks with distributed servers gives savings due to bandwidth reduction. Savings vary with number of users per session as well as with user distribution. For 3 users per session savings are approximately 40% and 35% for loosely/moderately concentrated users, respectively. For 6 users per session (more expensive), savings are in the order of 50% and 40% for loosely/moderately and highly concentrated distributions, respectively. Savings for highly concentrated users are lower than for loosely/moderately concentrated distributions because of its lower bandwidth requirements.

VI.2) Caching

The cost for cache replication can be easily generalized by the cost expression introduced in section V.2.We just need to consider that the demand of a cache located at level l is equal to U/d^l and the bandwidth cost per user is equal to $L(d) - l$. Figure 9 shows the total cache cost for networks with distributed server as a function of the normalized depth. We notice that the cache cost curve has a minimum at normalized depth of 0.7 and it has a similar shape to the shape of server replication curves. Moreover, we observe that for this type of networks caching is *not* worth adopting, when compared to server replication. We observe that savings due to the use of distributed servers (when compared to networks without stream sharing) can be as high as 50% for loosely/moderately and as high as 40% for highly concentrated users.

VI.1) A Brief Comparison

In order to understand the benefits of introducing distributed servers, we define C as the ratio between the cost of a distributed server and the cost of a standard server (used in networks without stream sharing). In figure 10, we plot the network cost for server replication for networks without stream sharing and the cost of networks with distributed server. We consider that in a real case the cost of a distributed server will be at most twice the cost of a standard server ($C = 2$). We notice that for loosely and moderately concentrated distributions we can have huge savings at almost any level of the distribution tree. For highly concentrated users, distributed servers is not worth adopting at levels higher than the optimum level, and attractive savings can be obtained if distributed server and standard server have the same cost. This is because the bandwidth cost for highly concentrated distributions may be quite lower than the bandwidth cost of loosely and moderately concentrated distributions. In Figure 11 we plot the cache replication cost for a cache size of 2000 programs. The trend of server replication cost is carried over to the cache replication cost, irrespective of the cache size.

VII) Conclusions

Video service is both a major business driver and a bandwidth consumer for the future broadband integrated network. Understanding video applications requirements is of paramount importance for network design. In this paper, we explore the interplay between bandwidth consume and program replication for different networks in providing Distributed Home Theatre services. We analyze server and cache replication strategies which aim at reducing high bandwidth demands. We also considered networks with both DHT and VoD services. We notice that design of networks with distinct capabilities may significantly differ. When analyzing the introduction of stream sharing, we notice that the number of users per session may impact the network cost due to the reduction of bandwidth cost. Savings on the network cost due to the adoption of stream sharing can be as high as 50% and it varies with the degree of sharing, i.e., savings are higher for highly concentrated users than for loosely/moderately concentrated users. In networks with both DHT and VoD services, the benefits of stream sharing are more striking for networks with higher percentage of DHT users (in which we have a higher degree of sharing).

Adopting distributed servers for the support of DHT is a very attractive solution. When compared to the cost of networks *without* stream sharing savings can be as high as 50%. Savings increases as we increase the number of users per session due to significant bandwidth reduction.

Network design for the provision of DHT services is highly influenced by the network capabilities. The design differences are mainly due to different degrees of bandwidth reduction For instance, the optimum placement level for server replication in networks *without* stream sharing and loosely/moderately concentrated users is around the normalized level 0.3 whereas for networks with distributed servers it is around level 0.7. Users distribution and the number of users per session also play a key role in network design. For example, the optimum placement level for server replication in networks with stream sharing is around the root of the distribution tree for loosely/moderately concentrated users while it is at level 0.4 for highly concentrated users. In general, server replication is worth adopting irrespective of the network characteristics. Furthermore, the cost of server replication is usually lower than is the cost of caching. Therefore, caching is not an interesting strategy.

Real cost will dictate the design decisions in a real network design. An attempt to understand findings for real network design was carried out. We varied the cost ratio between a distributed server and a standard server. We concluded that although employing distributed servers is worth adopting for loosely and moderately concentrated users, it is not always true for highly concentrated users.

An alternate approach which seems to be very attractive is the combination of stream sharing switching capability with distributed servers. These networks are the subject of our current investigation.

VIII) References

[1] J. A. Schnepf, D. H. C. Du, E. Russell Ritenour and A. J. Fahrmann, Building Future: Medical education environments over ATM networks”, *Communications of the ACM*, vol 38, n 2, pp 54-69, vol 38, n 2, Feb 1995.

[2] E.R. Madeira and N.L.S. Fonseca, Software Architecture for Multimedia Communications and Management, in Multimedia Communications Networks, Tatipamula and Bhumip (Ed), Artech House, 1997.

[3] *Statistical Abstract of United States*, U.S. Bureau of the Census, 113th ed., 1993.

[4] J. P. Nussbaummer, B. V. Patel, F. Schaffa and J.P.G. Sterbenz, “Networking requirements for interactive video on demand”, *IEEE Journal on Selected Areas in Communications*, vol 13, n 5, pp 779-787, June 1995.

[5] N.L.S. Fonseca, C.M.R. Franco and F. Schaffa, Network Design for the provision of Distributed Home Theatre, to appear in *Proc. of IEEE ICC’97*.

[6] *IEEE JSAC* special issue on Integrated Video-based Services for the Home, vol 14, n 6, August, 1996.

[7] J.P. Nussbaumer anf F. Schaffa, “Capacity analysis of CATV for on-demand multimedia distribution” in *Proc First ISMM Int Conf Distributed Multimedia Systems Applications*, 1994.

[8] J.P. Nussbaumer and F. Schaffa, “On Bandwidth and storage trade-off in multimedia distribution networks”, *Proc of INFOCOM’95*, pp. 1020-1026, 1995.

[9] Y. S. Chen, “Mathematical modeling of empirical laws in computer application: a case study”, *Comp. Math. Applicat*., pp. 77-87, Oct. 1992.

Acknowledgements

We would like to thank CNPq and FAPESP for their financial support and CENAPAD/SP for their computational support.

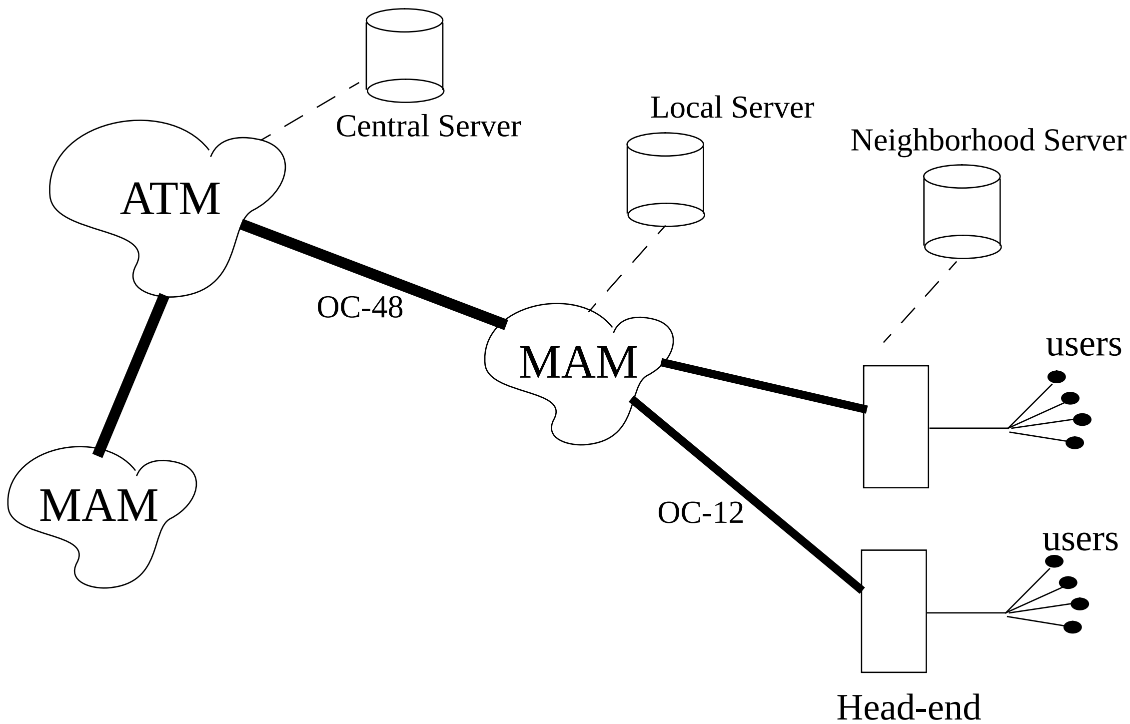

Figure 1: Network model.

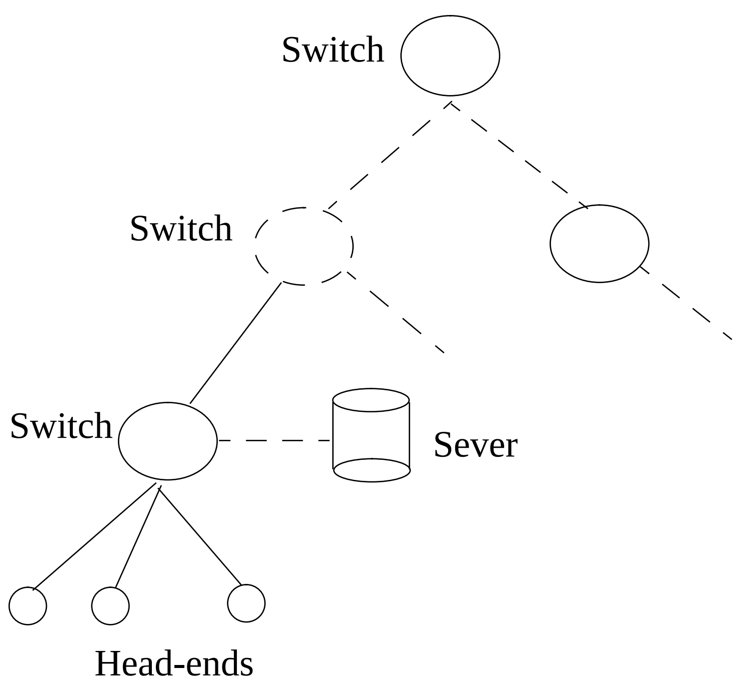

Figure 2: The Distribution tree.

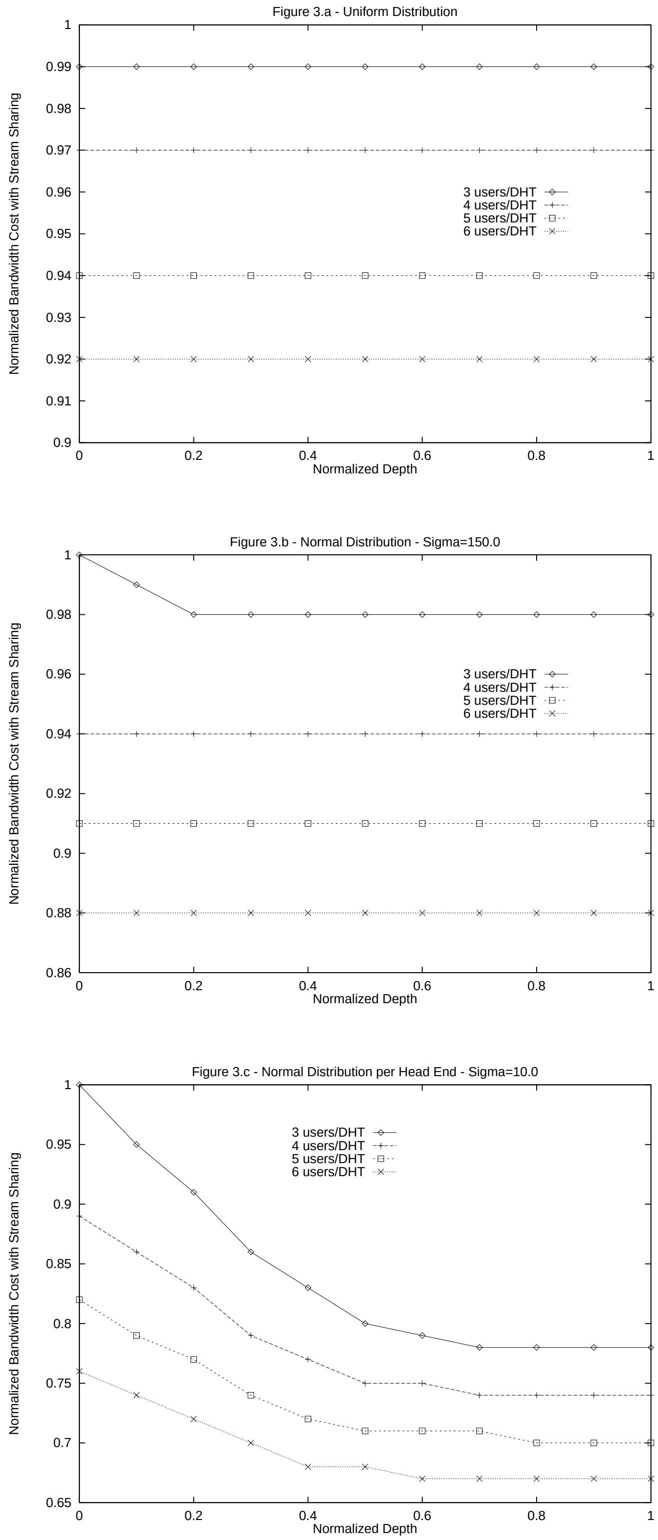

Figure 3: Normalized bandwidth cost of networks with stream sharing *x* normalized depth for: i) different number of users per DHT session, ii) a uniform distribution (Figure 3.a), iii) a normal distribution with σ = 150 (Figure 3.b) and iv) a normal distribution per head end with σ = 10 (Figure 3.c).

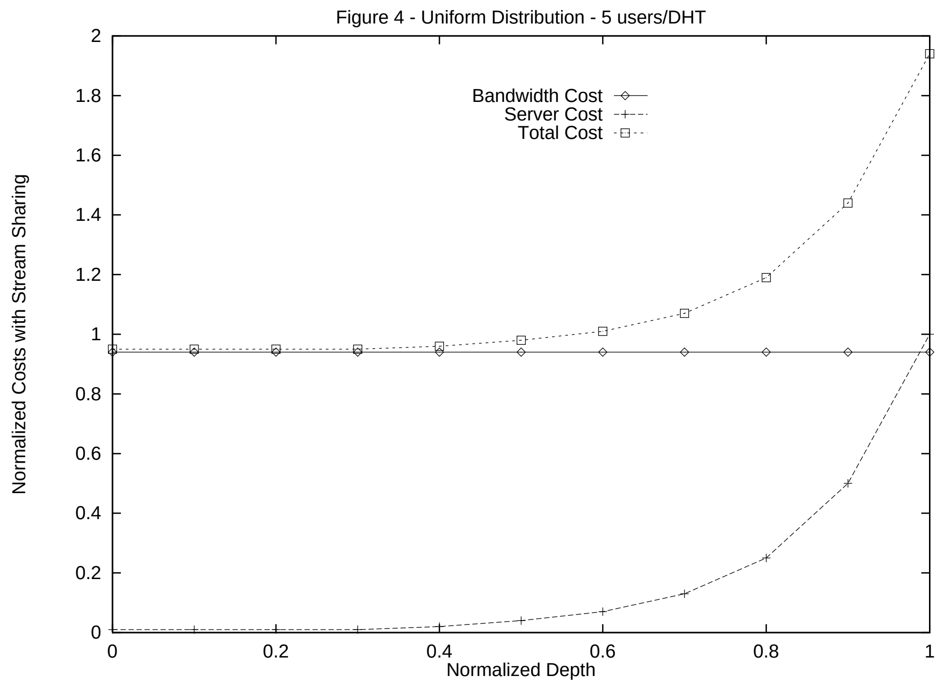

Figure 4: An example of total normalized cost computation considering server replication for a uniform distribution with 5 users per DHT session.

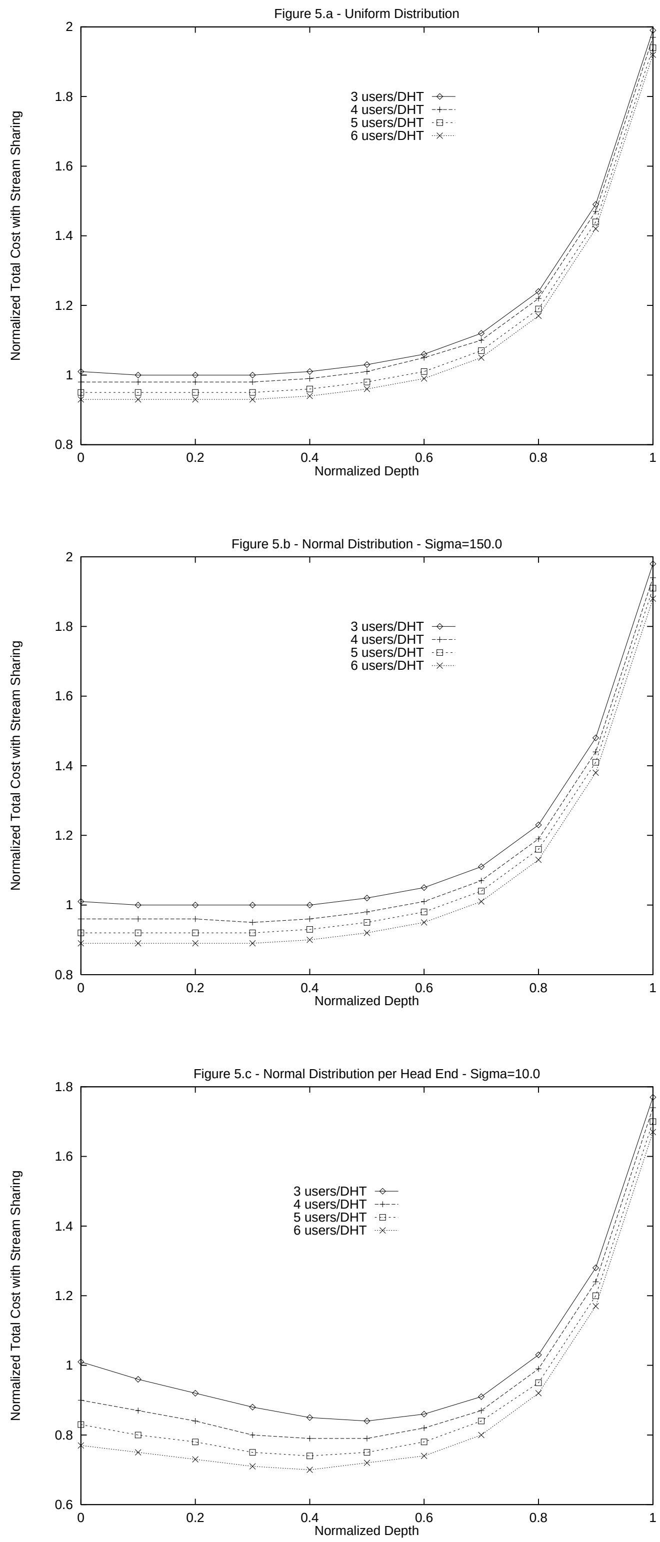

Figure 5: Normalized network cost of networks with stream sharing *x* normalized depth for: i) different number of users per DHT session, ii) a uniform distribution (Figure 5.a), iii) a normal distribution with $\sigma = 150$ (Figure 5.b) and iv) a normal distribution per head end with $\sigma = 10$ (Figure 5.c).

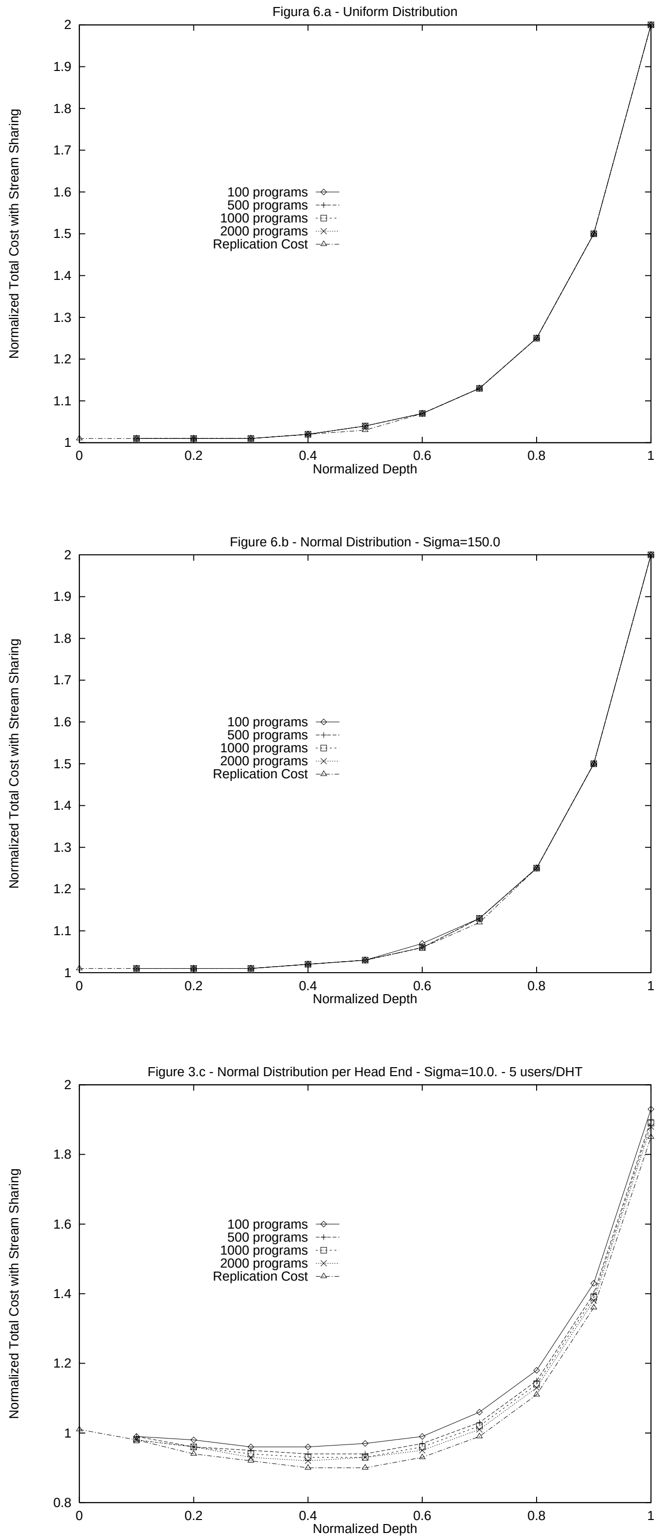

Figure 6: Normalized cache cost of networks with stream sharing *x* normalized depth for: i) different cache sizes, ii) 5 participants per DHT session, iii) a uniform distribution (Figure 8.a), iv) a normal distribution with $\sigma = 110$ (Figure 8.b) and v) a normal distribution per head end with $\sigma = 5$ (Figure 8.c).

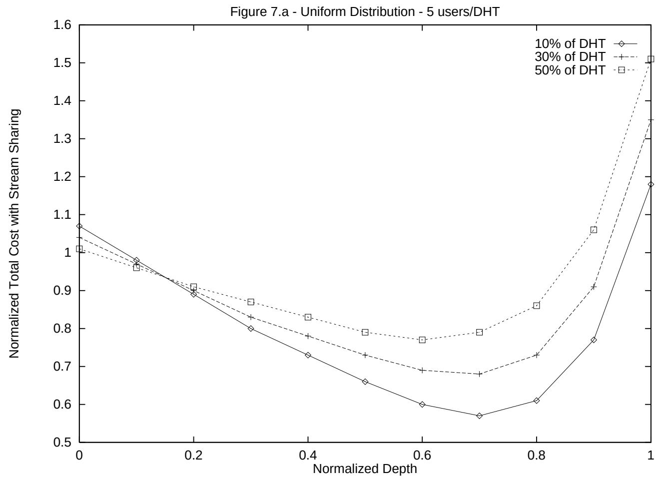

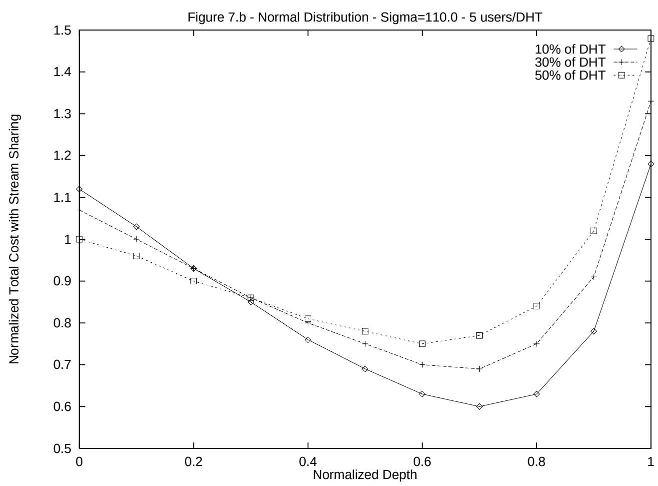

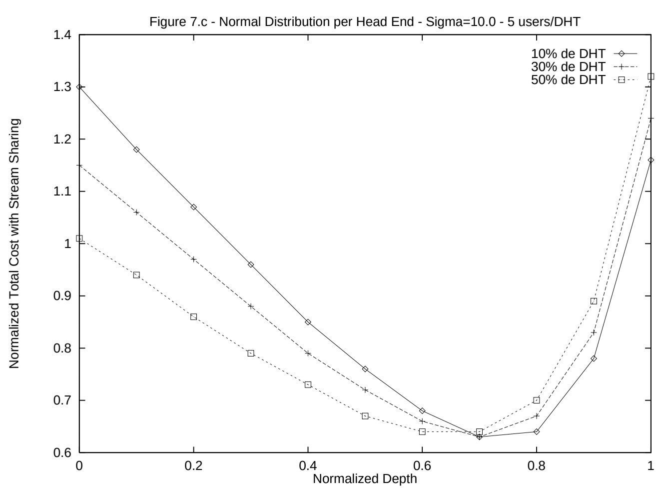

Figure 7: Normalized network cost of networks with stream sharing *x* normalized depth considering server replication in a network with both DHT and VoD services for: i) 5 users per DHT session, ii) uniform distribution (Figure 9.a), iii) a normal distribution with σ = 110 (Figure 9.b) and iv) a normal distribution per head end with σ = 5 (Figure 9.c).

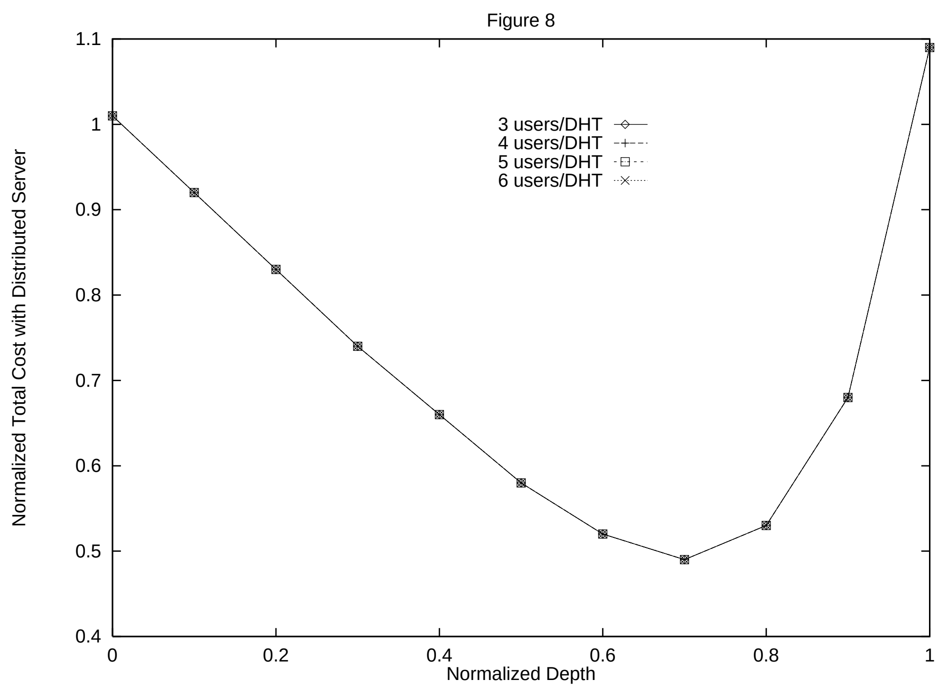

Figure 8: Normalized network cost of server replication of networks with distributed servers *x* normalized depth

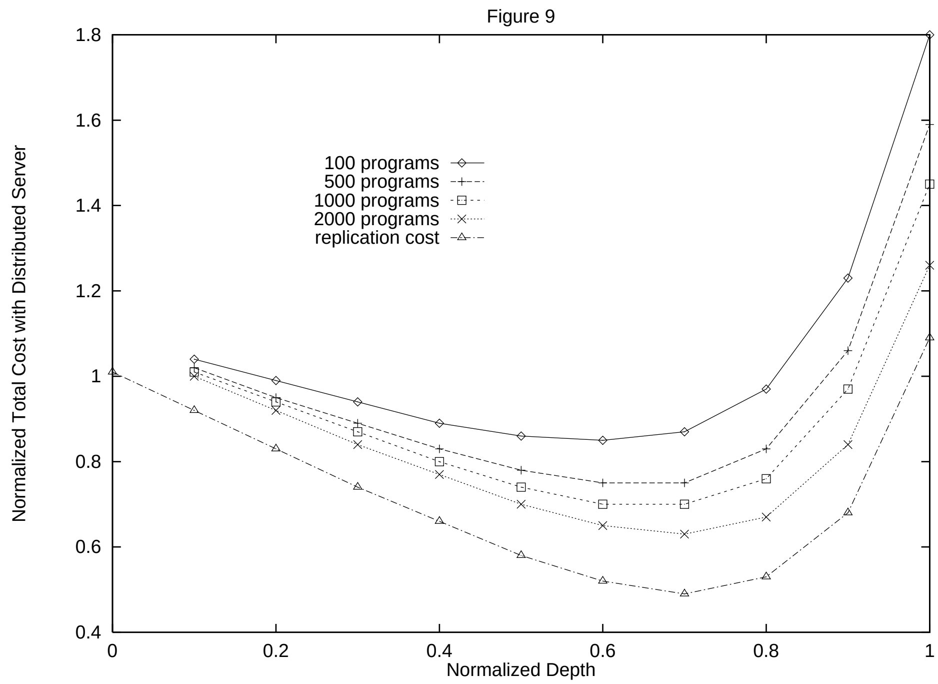

Figure 9: Normalized cache cost replication of networks with distributed servers *x* normalized depth for different cache sizes

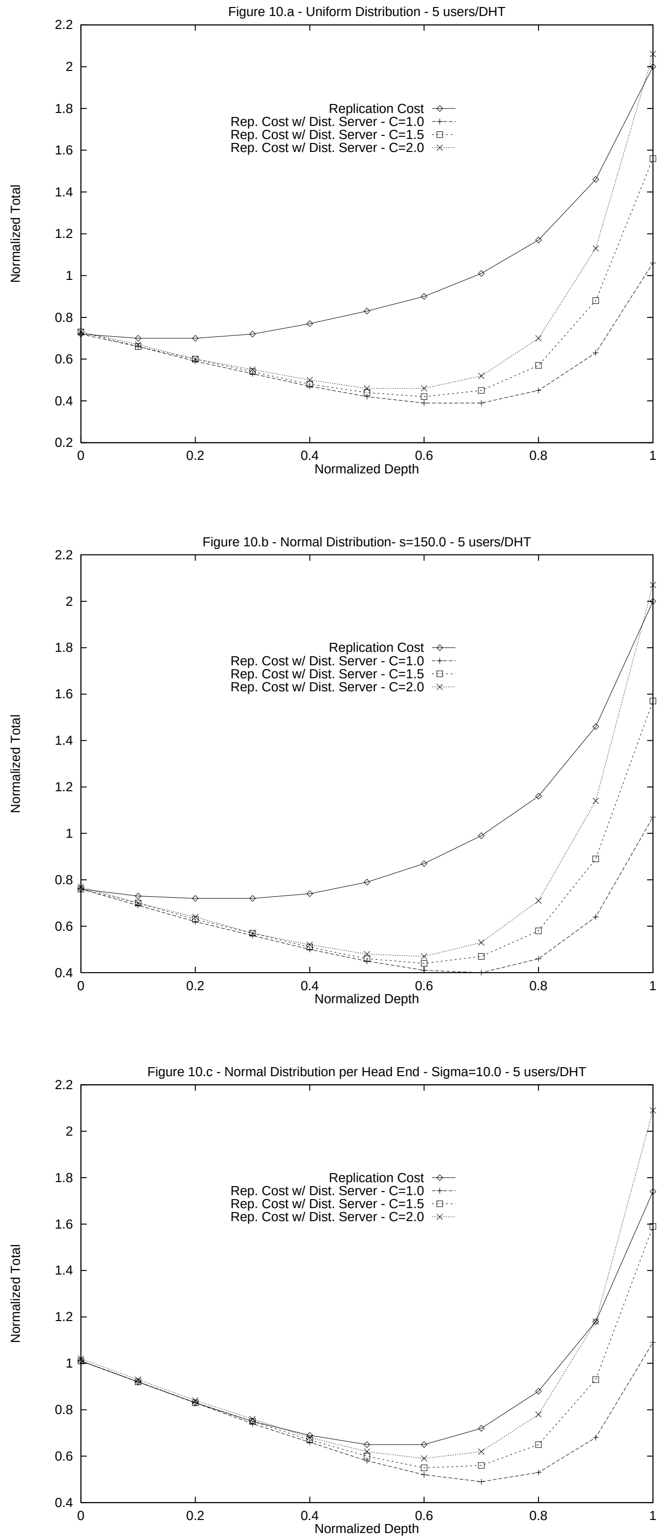

Figure 10: Normalized network cost for server replication *x* normalized depth for different values of distributed servers

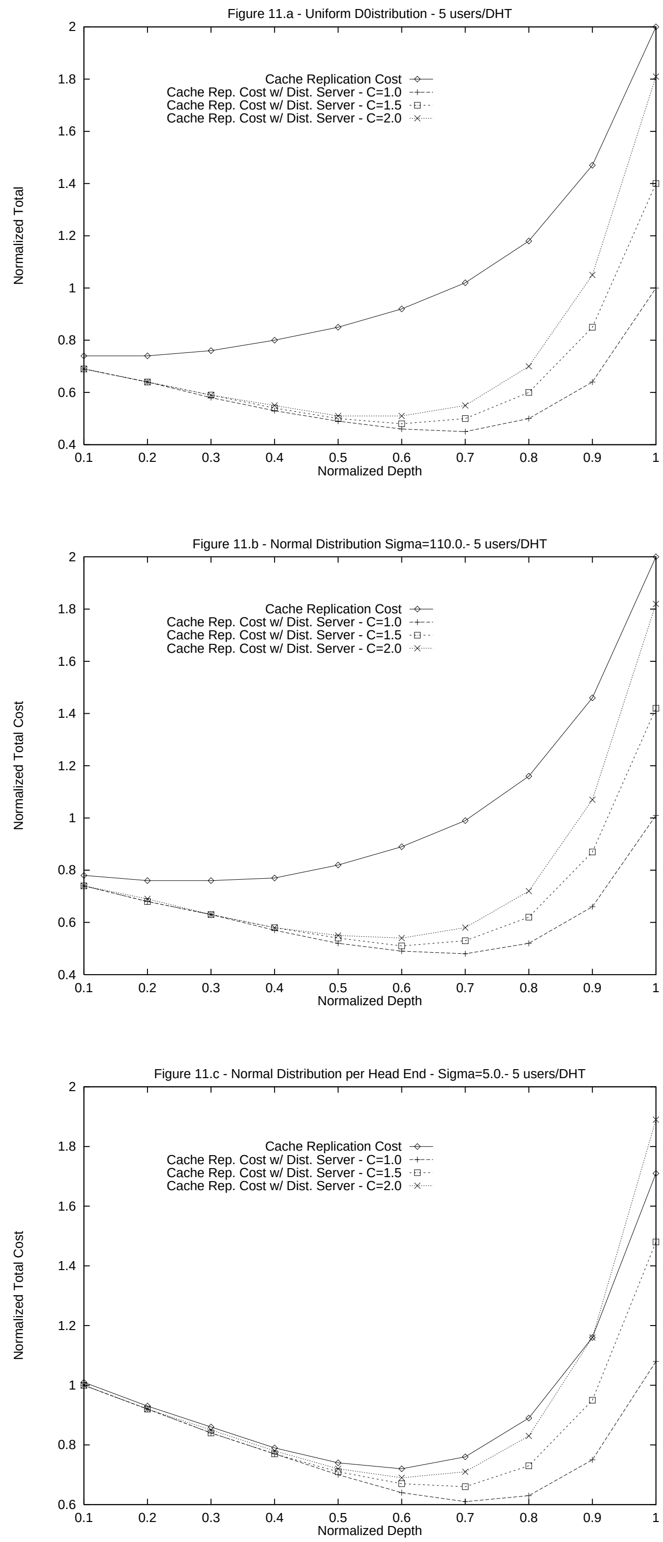

Figure 11: Normalized network cost for cache replication *x* normalized depth for different values of distributed servers