Using Multicast or a Combination of Unicast and Broadcast for Transmitting Popular Content

Jouni Karvo and Samuli Aalto
Helsinki University of Technology, Networking Laboratory
P.O.Box 3000, 02015 HUT
{jouni.karvo|samuli.aalto}@hut.fi

Abstract

We present a large deviations approximation for evaluating the gain of using multicast over unicast in a communications link. Also, we show how a combination of unicast and broadcast can be used to reduce the realised blocking probability, or to increase the number of supported users in a link.

1 Introduction

Traditionally, broadcast transmission has been used for transmitting popular content, such as for TV and radio transmission, where users can use their receivers to tune to the desired channel. For peer-to-peer communication, the unicast model has been used. Multicast is the method for group communication, where several users may join and listen the same transmission in a communications network.

The use of multicast has been advocated with the argument that it reduces the required network capacity by reducing the number of hop-by-hop transmissions, the gain of which has been studied in [7]. Another argument in favour of multicast is that it reduces the capacity required in a communications link, facilitating the increase in the number of users able to use the link with the target grade of service. The resulting gain has been estimated using a simple circuit-switched model [1]. The idea is to approximate the gain by calculating a supported number of users with a target blocking probability. For a small capacity link, the Engset formula gives the blocking probability of unicast connections, and the blocking probability of multicast connections can be estimated using Monte-Carlo simulation.

For larger links, however, the Monte-Carlo approach for estimating the multicast blocking probabilities has problems. This paper presents an approximative approach which is based on the theory of large deviations, notably on Hui's [4] work, and [8].

Martikainen et al. [6] suggest that for mobile devices, it would be beneficial to use broadcast for carrying the most popular content to the users, and to use multicast for the remaining content. There are many practical problems associated with implementing the multicast functionality in the network, see e.g. [3]. To overcome these problems, a combined approach can be used, where the most popular content is carried by broadcast but the remaining content is carried by unicast connections. We study this scenario, and note that the resulting use of channels can be improved over the use of multicast, provided that the operator is able to select the most popular channels for broadcasting.

This paper is organised as follows. Section 2 presents the multicast gain concept, referring to Aaltonen et al. [1]. Section 3 develops the large deviations approximation for multicast gain. A simple approach using the means of occupancy distribution is shown in Section 4, and the combination of unicast and broadcast is studied in Section 5. Finally, Section 6 summarizes the paper.

2 Multicast Gain

Aaltonen et al. [1] present a concept called "multicast gain". It is based on the study of a single link with capacity of C connections, which can either be the radio interface in a 3G network, or a link in the fixed network. The link carries traffic that consists of connections from users to specific sources, which differ in popularity. The term *channel* is used to draw an analogy to TV or radio broadcasts. The popularity of different channels is described by the *preference distribution.* Studies have shown that the Zipf distribution

$$a_i = \frac{i^{-\alpha}}{\zeta(\alpha)}, \qquad i = 1, 2, \ldots, \infty,$$

where $\zeta(\alpha)$ denotes the Riemann Zeta function, depicts human behaviour in this respect well.

For the most popular channels (small channel index i), capacity can be re-used by multicast or broadcast. Aaltonen et al. [1] present an estimation method for call blocking B^+_{mc} for the multicast case. For the broadcast case, the blocking probability for the channels that are transmitted is naturally zero. If channels are transmitted by unicast, the traditional Engset formula gives the call blocking probability B_{uc}.

Let us reverse the idea of call blocking probability calculation to a more dimensioning-like idea. Instead of the question "What is the grade of service with this number of users?" we ask "What is the number of users that can be served with the required grade of service?" The number of users that can be served can be achieved by fixing the activity level of individual users, and specifying the target blocking probability. Then iteration can be used to find the number of supported users for multicast N_{mc} and for unicast N_{uc}.

Now, the multicast gain η can be expressed as

$$\eta = \frac{N_{mc}}{N_{uc}}.$$

Naturally, a similar gain figure can be calculated for the combination of unicast and broadcast connections, after finding the supported number of users N_{ub}.

3 Large Deviations Approximation

This section develops a large deviations approximation for multicast gain, based on the Bahadur-Rao approximation [2, 4, 8]. The model is based on the assumption that the users are independent, and that there are so many of them that the traffic they generate follows approximately a Poisson process.

3.1 Unicast Blocking Probabilities

Let N_{uc} denote the number of unicast users, and p the probability that a user is active, i.e. has a connection. Then the traffic intensity of the corresponding Poisson process, ρ, is as follows:

$$\rho = pN_{uc}.$$

Now, let X denote the random variable describing the link occupation. The logarithmic moment generating function $\varphi(\beta)$ of this random variable,

$$\varphi(\beta) = \rho(e^\beta - 1).$$

To find the value of β that yields the Chernoff bound of $\mathrm{P}\{X \geq x\}$, let $\varphi'(\beta_x) = x$. Then,

$$\beta_x = \ln\frac{x}{\rho}.$$

The rate function $I(x)$ is as follows,

$$I(x) = \beta_x x - \varphi(\beta_x) = x\ln\frac{x}{\rho} - x + \rho.$$

An approximative value of $\mathrm{P}\{X \geq x\}$ (see [8, page 383]),

$$\mathrm{P}\{X \geq x\} \approx \frac{e^{-I(x)}}{\sqrt{2\pi}\beta_x\sigma(\beta_x)} = \frac{e^{-x\ln\frac{x}{\rho}+x-\rho}}{\sqrt{2\pi}\ln\frac{x}{\rho}\sqrt{x}}.$$

We use this approximation as the blocking probability approximation, i.e. $B_{uc} \approx \mathrm{P}\{X \geq C\}$.

3.2 Multicast Blocking Probabilities

The multicast model consists of a preference distribution, where the popularity of a multicast channel is determined by two quantities; the probability that a user is active, p, and the probability that a user selects channel i, a_i. We consider the Zipf distribution:

$$a_i = i^{-\alpha}\zeta(\alpha)^{-1}, \qquad i = 1, \ldots, \infty,$$

where α is the distribution parameter and $\zeta(\alpha)$ is the Riemann Zeta function with parameter α, i.e.

$$\zeta(\alpha) = \sum_{i=1}^{\infty} i^{-\alpha}.$$

There is an inherent dependency between channels in the multicast scenario, since one user is assumed to have only one connection at a time. When the capacity of the link increases, however, the number of users increases quickly. Thus, we assume that this dependency can be neglected so that each channel can be modelled using a Bernoulli variable, with the probability p_i that channel i is on, (as the probability of an M/M/∞ queue being non-idle [5])

$$p_i = 1 - e^{-pN_{mc}a_i}.$$

Since the channels are assumed independent, the moment generating function of the link occupancy distribution is a product of the moment generating functions of these Bernoulli variables. Thus, the logarithmic moment generating function of link occupancy X,

$$\varphi(\beta) = \sum_{i=1}^{K} \ln(1 - p_i + p_i e^\beta).$$

The parameter β_x is calculated as for the unicast case, but in this case, a numeric approach is required.

For $K = \infty$, an additional problem arises, requiring a new approximation by truncating the distribution and using another approximation for the tail. We approximate the effect of the tail by first noting that for the channels in the tail, the probability that there is more than one user connected to the channel is very small. Thus, effectively these channels are carried as unicast on the link. Let ρ_t denote the traffic intensity for the channels in the tail,

$$\rho_t = \frac{pN_{mc}\zeta(\alpha, K+1)}{\zeta(\alpha)},$$

where K is the number of the last channel to be treated with the first approximation, and

$$\zeta(\alpha, x) = \sum_{i=x}^{\infty} i^{-\alpha}, \qquad \zeta(\alpha, 1) = \zeta(\alpha).$$

Now, the logarithmic moment generating function

$$\varphi_K(\beta) = \sum_{i=1}^{K} \ln(1 - p_i + p_i e^{\beta}) + \rho_t(e^{\beta} - 1).$$

Then,

$$\varphi_K'(\beta) = \sum_{i=1}^{K} \frac{p_i e^{\beta}}{1 - p_i + p_i e^{\beta}} + \rho_t e^{\beta}.$$

$$\varphi_K'(\beta_x) = x.$$

$$\varphi_K''(\beta) = \sum_{i=1}^{K} \frac{(1-p_i)p_i e^{\beta}}{(1 - p_i + p_i e^{\beta})^2} + \rho_t e^{\beta}.$$

$$\sigma(\beta_x) = \sqrt{\varphi_K''(\beta_x)}.$$

$$I(x) = \beta_x x - \varphi_K(\beta_x).$$

An approximative value of $\mathrm{P}\{X \geq x\}$,

$$\mathrm{P}\{X \geq x\} \approx \frac{e^{-I(x)}}{\sqrt{2\pi}\beta_x \sigma(\beta_x)}.$$

This approximation is used as the approximation of multicast blocking probabilities, $B_{mc} \approx \mathrm{P}\{X \geq C\}$.

3.3 Numerical results

Figure 1 shows some numerical results as a function of α. For example, if $\eta = 10$ multicast gain were expected, the α required for links of 100 channels capacity would be about 1.7, but for links of 500 channels, $\alpha \approx 1.5$ would suffice.

The problem with the large deviations approach is the selection of K. The error introduced in the approximation depends heavily on selection of K. The error in the multicast part of the approximation underestimates the blocking probability while the error for the unicast part of the estimation overestimates the blocking probability.

4 Comparison between means

It is also possible to estimate the multicast gain by comparing the means of link occupancy distributions, instead of call blocking probabilities. This section develops such an estimation.

The estimation is based on calculating the number of users with activity level p that generate mean link occupancy C, both for unicast, where

$$N_{uc} = C/p,$$

and for multicast. The number of multicast users that can be supported is solved numerically from the following equation:

$$\sum_{i=1}^{\infty}(1 - e^{-pN_{mc}i^{-\alpha}/\zeta(\alpha)}) = C.$$

The mean-value version of multicast gain is then calculated by dividing N_{mc} by N_{uc}.

5 Combination of Unicast and Broadcast

The idea in the combination of unicast and broadcast is to use broadcast to transmit the most popular content, and to let users use the normal unicast mode for the not so popular content. To make this method work, the operator needs to select the appropriate content to be broadcast in the network, and to decide what is the correct number K_y of the most popular content that deserve a broadcast channel.

The blocking probability for this scenario is easily calculated. There is no blocking for the broadcast channels, and for the remaining channels the standard Engset formula can be used.

First, for the Zipf preference distribution, the tail weight (the fraction of the calls that use unicast) is calculated as

$$w = \frac{\zeta(\alpha, K_y + 1)}{\zeta(\alpha)}.$$

Then, the call blocking probability for n users will be

$$B_{ub} = w\frac{\binom{n-1}{C-K_y}(pw)^{C-K_y}(1-pw)^{n-1-C+K_y}}{\sum_{i=0}^{C-K_y}\binom{n-1}{i}(pw)^i(1-pw)^{n-1-i}}.$$

Figure 2 shows some numerical results. The selection of K_y affects the realised blocking probability. Knowing the preference distribution, it is possible to choose K_y such that it minimises the blocking probability. Note that for a wide range of K_y, the resulting blocking probability is smaller than for multicast. The value for $K_y = 0$ yields the unicast blocking probability, and $K_y = C$ yields the broadcast blocking

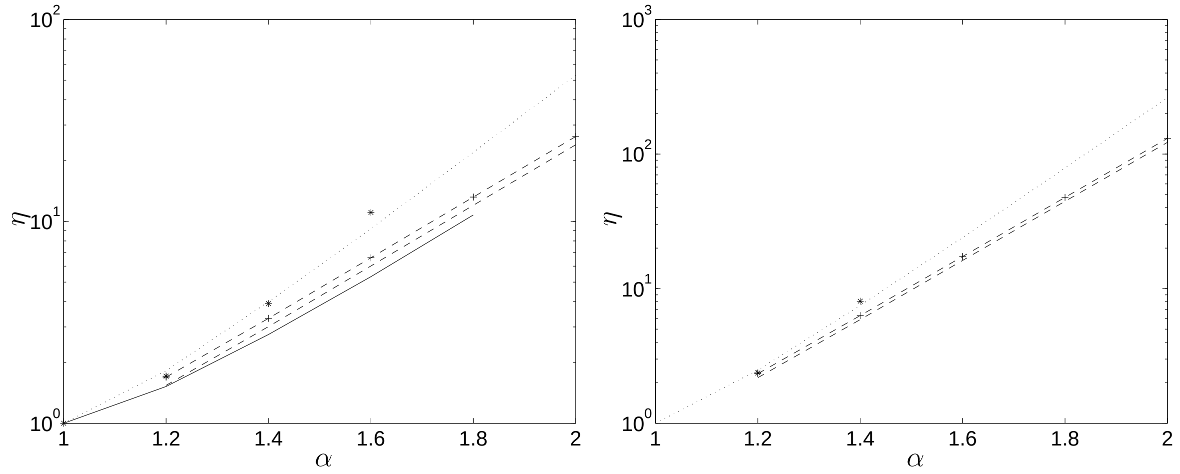

(a) $C = 100$, $K = 20$, $p = 0.1$ (b) $C = 500$, $K = 100$, $p = 0.1$

Figure 1: multicast gain vs. α. Solid line: MC simulation, dashed line: the large deviations approximation for multicast, unicast from the Engset formula. Dashed line with pluses: large deviations approach for both unicast and multicast. Dotted line: comparison between means. Asterisks: combination of unicast and broadcast.

probability (which equals the tail weight w, and is independent of the number of users in the system).

Figure 1 shows some results on comparing multicast and a combination of unicast and multicast with optimal K_y adjustment. For large values of α, the combination of unicast and broadcast outperforms use of multicast. This property is even so strong that for large enough links (large C), and big α, the number of users of the link do not affect the blocking probability at all. See Figure 3 for the capacity required for a link to accommodate an infinite number of users with 2% target blocking probability, as a function of α.

6 Summary and further work

This paper presents a large deviations approximation for estimating the blocking probabilities for multicast streams, and uses that approximation for estimating multicast gain for large links, where Monte Carlo simulation has problems. We also show a simpler estimate for multicast gain, namely a method for comparing the user populations generating meanwise similar link occupations. The results of both of these estimations produce similar results.

Instead of using pure unicast or multicast for transmitting content, it is possible to use a combination of unicast and broadcast, assuming the most popular content to be known to the operator. We show how the blocking probability for such a scenario is calculated, and compare the gain of this scenario to the multicast gain. We also note that for large enough links, use of pure broadcast allows an infinite number of users to be served with the specified target blocking probability.

As further work, the large deviations blocking probability approximation could be based on the truncated distributions instead of the tail probability estimation. Also, the broadcast-unicast scenario could be studied so that the criterion would not be the call blocking probability but instead the channel blocking probability of the least used channel.

References

[1] Aaltonen, J., Karvo, J., and Aalto, S. Multicasting vs. unicasting in mobile communication systems. In *Proc. Workshop on Wireless Mobile Multimedia, WoWMoM 2002* (Atlanta, GA, Sept. 2002), pp. 104–108.

[2] Bahadur, R. R., and Rao, R. R. On deviations of the sample mean. *Annals of Mathematical Statistics* (1960), 1015–1027.

[3] Diot, C., Dabbous, W., and Crowcroft, J. Multipoint communication: A survey of protocols, functions, and mechanisms. *IEEE Journal on Selected Areas in Communications 15*, 3 (Apr. 1997), 277–290.

[4] Hui, J. Y. Resource allocation for broadband networks. *IEEE Journal on Selected Areas in Communications 6*, 9 (Dec. 1988), 1598–1608.

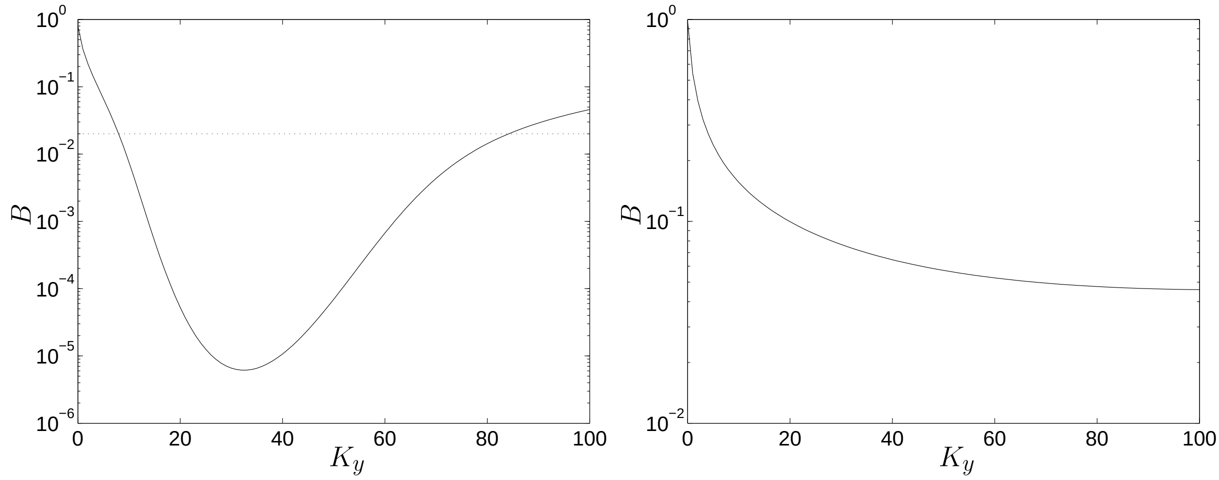

(a) $n = 4700$, $p = 0.1$. Solid line: blocking probability for the combination of unicast and broadcast channels, dotted line: blocking probability for multicast ($B_{mc} = 0.02$).

(b) $n = 40700$, $p = 0.1$. Solid line: blocking probability for the combination of unicast and broadcast channels.

Figure 2: Combination of unicast and broadcast. $\alpha = 1.6$, $C = 100$.

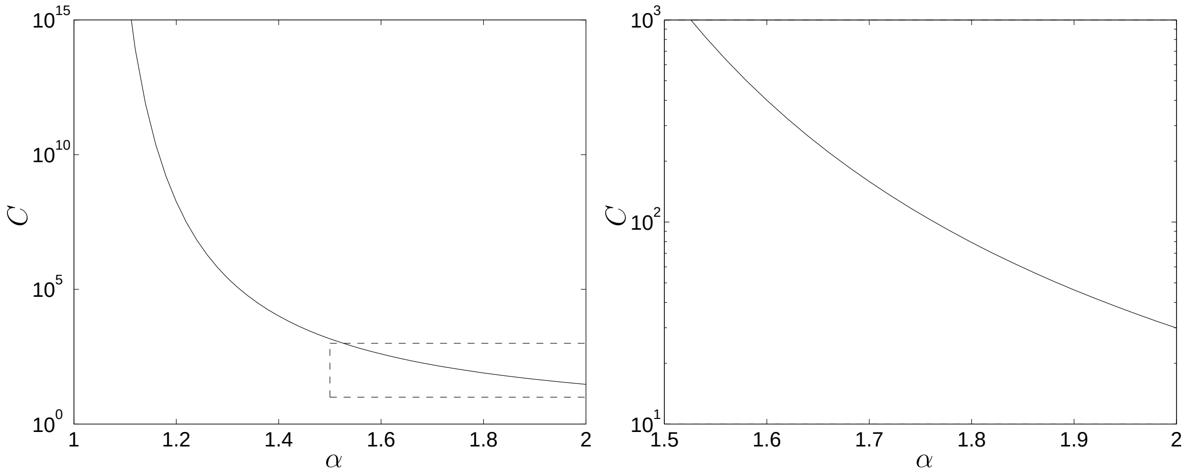

Figure 3: The number of broadcast channels required to achieve a 2% blocking probability for a Zipf preference distribution with parameter α. The figure on the right shows the dashed rectangle in a greater detail.

[5] Karvo, J., Virtamo, J., Aalto, S., and Martikainen, O. Blocking of dynamic multicast connections in a single link. In *Proc. of International Broadband Communications Conference, Future of Telecommunications* (Stuttgart, Germany, Apr. 1998), pp. 473–483.

[6] Martikainen, O., Nyman, E., Räsänen, J., and Koponen, P. Terminal mobility in broadband networks. In *Proc. 3rd IFIP TC6 Workshop on Personal Wireless Communications* (Tokyo, Japan, Apr. 1998), pp. 239–246.

[7] Mieghem, P. V., Hooghiemstra, G., and van der Hofstad, R. On the efficiency of multicast. *IEEE/ACM Transactions on Networking 9*, 6 (Dec. 2001), 719–732.

[8] Roberts, J., Mocci, U., and Virtamo, J., Eds. *Broadband Network Teletraffic.* COST 242, Final Report of Action. Springer, 1996.