

Access network dimensioning for elastic traffic based on flow-level QoS

Pasi Lassila and Jorma Virtamo

Networking Laboratory, Helsinki University of Technology,
P.O. Box 3000, FIN 02015 HUT, Finland

Email: {Pasi.Lassila, Jorma.Virtamo}@hut.fi

June 2, 2005

Abstract

Abstract: Network dimensioning for circuit switched networks has long been a well understood problem. However, dimensioning for data traffic, which is by nature elastic, is still an area that is not very well developed. Some work has been recently published addressing data network dimensioning and the present paper presents a survey on these. In our opinion, robustness of the used network models, forming the foundation of any dimensioning methods, is the key issue. Therefore, in this paper we propose the use of models based on the notion of balanced fairness. Dimensioning is important for example in access networks and these models allow simple and explicit evaluation of the performance of tree type access networks. We define the dimensioning problem as an optimization task where the link capacities are chosen such that costs are minimized subject to per-flow constraints on the minimum flow-level throughput. The approach is illustrated through various numerical examples.

1 Introduction

Optimal network dimensioning is a classical problem in teletraffic theory. For circuit switched networks the dimensioning criterion is naturally the call blocking probability, and the dimensioning methods have been based on the classical Erlang formula [10]. One attractive property of the Erlang formula is its insensitivity property, i.e., the performance of the system only depends on the load and not on any detailed properties of call holding times. This insensitivity property partly explains why the Erlang formula is still used for

dimensioning despite that the traffic characteristics have changed considerably since the publication of the result in 1917.

Here our focus is on dimensioning of data networks, such as IP networks, which carry mostly data traffic. The traffic is by nature elastic, which implies that the applications transmitting the data can handle temporary fluctuations in the transmission rate, as is the case with applications sending data over TCP or some other protocol with similar rate control behavior (i.e., a TCP-friendly protocol). For dimensioning purposes, the network models must be simple enough to allow evaluation of the model for realistic size networks. Hence, detailed models taking into account too much details of the used network protocols, for example properties of TCP, can not be used. In addition to simplicity, robustness of the used network models is another key issue. From the point of view of dimensioning elastic traffic the models necessarily need to take into account at least the way in which the flows share the bandwidth. Common idealized bandwidth sharing schemes are max-min fairness [1], proportional fairness [11], and more recently, balanced fairness [3]. In practice, none of the above is exactly realized by TCP. Thus, none of them can be considered as the “correct one”, and for dimensioning purposes the one giving the most practical models should be considered.

The recently introduced notion of balanced fairness provides an elegant abstraction of the way that the resources of the network are shared by competing flows in a dynamic situation, where the flows arrive and leave the system stochastically. Network models based on balanced fairness have several attractive features (for proofs, see, e.g., [3]). These models are insensitive, similarly as the model used in Erlang’s formula, and hence the models are inherently robust. Furthermore, balanced fairness is also the most efficient insensitive allocation. Additionally, the balanced fair allocation approximates well other allocations such as max-min fairness (see [3] and [18]), but contrary to max-min fair allocations, the performance of balanced fair allocations can be computed explicitly for various network topologies. Also, max-min fairness or proportional fairness do not, in general, yield insensitive allocations. Efficient recursive algorithms for solving the performance have been derived in [5] for a general network topology and for trees in [6].

Due to the aforementioned properties of balanced fairness, we propose in this paper the use of these models as appropriately robust tools for network dimensioning. However, our dimensioning approach is only to apply these detailed models at the point of the network where bandwidth can be considered a scarce resource. In fixed networks, this can be argued to be the case only in access networks having typically a tree structure. The dimensioning problem in tree-type access networks is formulated as an optimization problem, where the link capacities are determined such that the total cost of the network is minimized subject to a constraint on the minimum per-flow throughput. By using the so called store-and-forward bound (derived in [4]) as an approximation to the throughput, the optimization problem can be solved efficiently numerically resulting in conservative estimates of the required capacity. The amount of over-provisioning caused by the use of the store-and-forward bound can be verified against the exact solution under balanced fairness by using the recursive algorithm given in [6]. Alternatively, if greater accuracy in the solutions is required, one can use

the so called parking lot approximation. However, it does not yield a proven bound for the performance, but instead the results are approximative in nature (although typically rather accurate ones). In this paper, we also indicate how the parking lot approximation can be extended to include also the impact of per-flow access rate limitations. Finally, our proposed approach is illustrated through several numerical examples, where both the store-and-forward bound and the parking lot approximation have been used.

The paper is organized as follows. Related work is discussed in Section 2. Section 3 contains a brief exposition of balanced fairness, and our dimensioning approach is detailed in Section 4. Numerical examples are presented in Section 5 and conclusions are given in Section 6.

2 Related work

2.1 Models for network dimensioning

A general framework for dimensioning of modern packet switched networks has been outlined in [9]. According to the framework, traffic is classified into three classes: realtime services, interactive data and streaming services, and delay-tolerant services. For these classes, relevant QoS parameters to be used as basis for the dimensioning are identified (packet loss, blocking probability, throughput), and a four stage approach is outlined that describes the incremental process of dimensioning the network to support all three traffic classes. In conclusion, the focus in [9] is to develop a framework for performing capacity planning, and not to propose a specific solution to the entire problem.

In this paper, we are concerned with dimensioning the network for elastic traffic, i.e., according to [9], for delay-tolerant services, and for streaming and interactive data services. Existing approaches for this can be broadly assigned a few distinguishing properties. One important property concerns whether the model is for dimensioning an arbitrary network or just a single bottleneck link. Another distinction can be made according to whether the model treats the traffic as fluid or does the model employ some dynamic (stochastic) model of the traffic. First we discuss methods for dimensioning a general network and then briefly some recent approaches on bottleneck link dimensioning.

Existing methods in the literature on dimensioning of arbitrary data networks are usually formulated as optimization problems, but the distinction comes from the assumed network traffic model, i.e., whether it is static or dynamic. Static is here used to imply that the network model treats traffic basically as fluid and no stochastic elements are present in the model. The dimensioning problem is then formulated as a multi-commodity optimization problem, see [16, 12, 14, 13, 15]. The network consists of routes (paths) and the problem is to determine the link capacities and the bandwidth allocations for each route such that a given utility function is maximized. The maximization is then performed by assigning a budget constraint on the total cost of the network links. The utility function can be formulated such that overall throughput is maximized [13], a max-min fairness criteria is

maximized [13], or a proportional fairness criteria is maximized [16, 12, 14, 13]. The multi-commodity flow problem also easily allows additional considerations to be incorporated in the basic model, such as inclusion of network failures using backup paths [13] or bounded flows [14, 16]. In principle, the multi-commodity flow problem is well defined and easy to apply for various types of network design problems, as is comprehensively discussed in [15]. However, the approach does neglect one important aspect - the dynamic and stochastic nature of the traffic. The models do not include in any way the notion of offered traffic whose needs must be satisfied according to a given performance criteria.

A classic packet network dimensioning method that is also formulated as an optimization problem, but employing a stochastic traffic model, is the square-root method in [1]. The idea in the method is basically to determine the link capacities such that the total network cost is minimized subject to a constraint on the overall mean packet delay in an open M/M/1 queuing network. The drawback of this method is that the performance criterion is a packet level metric (mean queuing delay) that cannot easily be related to something meaningful end-to-end for a user. Also, the underlying network model makes strong assumptions on the traffic, i.e., it is not robust.

We propose the use of models based on balanced fairness for dimensioning of access networks. Our formulation is in spirit similar to the one in [1], except that the performance measure in our case is not the packet delay in a network of M/M/1 queues, but the flow level delay (or equivalently throughput) associated with file transfers in a network exhibiting balanced fairness. Another difference is that the formulation in [1] uses a single constraint on the overall mean packet delay, whereas in our case we use a separate constraint for each class. However, when using the store-and-forward bound and substituting the per class mean flow transfer delay requirements with an overall mean flow transfer delay of the network, the optimization problem, in fact, becomes equivalent to the square-root method in [1].

2.2 Models for link dimensioning

Finally, there are a number of papers that focus on dimensioning of a single bottleneck link. In [7], a number of different models have been used to evaluate the dimensioning requirements given by different single link traffic models. The results are also compared against measurements. Link provisioning is considered in [19], where the problem is to determine the amount of bandwidth needed over a relatively short time interval, say 5 minutes, such that the amount of incoming traffic does not exceed the link rate.

From the robustness point of view, models based on processor sharing (PS) queues are the most attractive ones for link dimensioning. For such models (and also for users themselves), a natural dimensioning criterion is the per flow throughput (or alternatively the file transfer delay). The so called GPS (generalized processor sharing) model [8] has been applied for dimensioning of elastic TCP traffic in [17]. The GPS model assumes an infinite user population (Poisson arrivals) but it can nicely capture the situation where flows are peak rate limited such that the bottleneck link appears as an M/G/ ∞ system until the link becomes full, after which the flows start sharing the bandwidth according to PS. A finite user

population variant of GPS has been studied in [2] where a simple and explicit formula has been derived for the required capacity given the number of users, offered traffic and a target dimensioning criteria for the mean useful rate of a user.

3 Short review of balanced fairness

The following description of balanced fairness in general topology networks has been adopted from [6]. The network consists of links $\mathcal{L} = \{1, \dots, L\}$, where link l has capacity C_l . A random number of flows compete for the bandwidth of these links. There are N classes of flows, $\mathcal{F} = \{1, \dots, N\}$, where each class i is characterized by a route \mathcal{R}_i consisting of a set of links. When link l is on route \mathcal{R}_i we use the notation $l \in \mathcal{R}_i$. Conversely, defining $\mathcal{F}_l \subset \mathcal{F}$ to be the set of flow classes going through link l we can write equivalently $i \in \mathcal{F}_l$. The mean volume of information offered by flows in class i per unit time, i.e., the load of class i , is denoted by ρ_i . The network state is defined by the vector $x = (x_1, \dots, x_N)$, where x_i is the number of class- i flows in progress.

3.1 Bandwidth allocation under balanced fairness

The total capacity $\phi_i(x)$ allocated to class- i flows is assumed to be shared equally between these flows and to depend on the network state x only. The capacity allocation must satisfy,

$$\sum_{i \in \mathcal{F}_l} \phi_i(x) \leq C_l, \quad \forall l \in \mathcal{L}. \quad (1)$$

The allocation is said to be balanced if

$$\frac{\phi_i(x - e_j)}{\phi_i(x)} = \frac{\phi_j(x - e_i)}{\phi_j(x)}, \quad \forall i, j, \quad x_i > 0, \quad x_j > 0,$$

where e_i is an N -vector with 1 in component i and zeros elsewhere. The balance property implies that there is a balance function $\Phi(x)$ such that

$$\phi_i(x) = \frac{\Phi_i(x - e_i)}{\Phi_i(x)}, \quad \forall i, \quad x_i > 0.$$

Basically, any positive function $\Phi(x)$ defines a balanced allocation. However, there is a unique balanced allocation such that for any network state x all the capacity constraints (1) are satisfied and at least one of them is satisfied as an equality, i.e., at least one link is saturated. For this allocation, called *balanced fairness*, the balance function $\Phi(x)$ is obtained recursively from

$$\Phi(x) = \max_l \left\{ \frac{1}{C_l} \sum_{i \in \mathcal{F}_l} \Phi(x - e_i) \right\}. \quad (2)$$

The model can be extended to take into account per-flow rate limitations due to, for example a limited access rate. Let a_i denote the rate limit of class i flows. The presence of rate limits results in additional constraints,

$$\phi_i(x) \leq x_i a_i, \quad \forall i, x_i > 0.$$

The balance function is then obtained recursively from

$$\Phi(x) = \max \left\{ \max_l \left\{ \frac{1}{C_l} \sum_{i \in \mathcal{F}_l} \Phi(x - e_i) \right\}, \max_{i: x_i > 0} \left\{ \frac{\Phi(x - e_i)}{x_i a_i} \right\} \right\}. \quad (3)$$

3.2 Throughput of the system

For networks exhibiting the balanced fairness property, the throughput of class- i flows, γ_i , is defined as $\gamma_i = \mathbb{E}[S_i]/\mathbb{E}[T_i]$, where $\mathbb{E}[S_i]$ is the mean flow size for class i and $\mathbb{E}[T_i]$ is the mean sojourn time for class i . Applying Little's result yields

$$\gamma_i = \frac{\rho_i}{\mathbb{E}[x_i]} = \frac{G}{\partial G / \partial \rho_i} = \frac{1}{\partial \log G / \partial \rho_i}, \quad (4)$$

where G is the normalization constant,

$$G = \sum_{x_1=0}^{\infty} \cdots \sum_{x_N=0}^{\infty} \Phi(x) \rho_1^{x_1} \cdots \rho_N^{x_N}.$$

For access networks, topologies of special interest are tree networks. For such networks it turns out that the allocation defined by (2) or (3) is Pareto efficient, and the recursion can be computed efficiently even for relatively large trees [6]. If the computation of the throughput using the algorithms that implement the recursions (2) or (3) is computationally too intensive, approximations to the flow throughput can also be used.

The store-and-forward bound, as derived in [4], yields a proven lower bound for the throughput of class- i flows in any network topology. It is simple and gives a conservative estimate of the throughput, which is potentially useful in practical applications, such as network dimensioning. To be exact, the store-and-forward approximation for the throughput of class- i flows, γ_i^{SF} , equals

$$\gamma_i^{\text{SF}} = \left(\frac{1}{a_i} + \sum_{l \in \mathcal{R}_i} \frac{1}{C_l - R_l} \right)^{-1}, \quad (5)$$

where $R_l = \sum_{i \in \mathcal{F}_l} \rho_i$ denotes the total load of link l . The interpretation of the lower bound is that transmission of files in the network occurs as if each node along the flow's route first waits until it has the entire file stored and then transmits it on the link to the next node. The store-and-forward approximation is typically accurate for class- i if there is a distinctive bottleneck along the route \mathcal{R}_i . If access rate limits do not need to be considered, the first term $1/a_i$ is simply removed.




Figure 1: The parking lot configuration.

If the store-and-forward bound is considered too inaccurate, it is also possible to utilize the so called parking lot approximation (i.e., it is not an upper or lower bound), which is typically very accurate for tree networks. The basic parking lot configuration without access rate limitations is illustrated in Figure 1. It consists of $L = n$ links with capacities $C_1 \leq C_2 \leq \dots \leq C_n$, which carry $N = n$ classes of flows. Class-1 flows go through all the links, class-2 flows go through links 2 to n , and class- i flows go through link i to n . For this configuration the throughput of each class can be solved explicitly [5],

$$\gamma_i = \left(\frac{1}{C_i - R_i} + \sum_{l=i+1}^n \left(\frac{1}{C_l - R_l} - \frac{1}{C_l - R_{l-1}} \right) \right)^{-1}, \quad (6)$$

where $R_i = \sum_{j=1}^i \rho_j$, i.e., the total load of link i . Note that this approximation includes the same terms as the earlier store-and-forward bound, where conceptually the transmitted files were first stored entirely at each node before being transmitted onward. However, when a file is being transmitted it is actually transmitted concurrently by all nodes along the route. The parking lot approximation improves the store-and-forward bound by including additional terms (those with the minus sign) that take into account the concurrency at least between two neighboring nodes.

The parking lot configuration can be used to approximately model the situation as observed by class i flows in a tree, including per-flow access rate limits a_i . The idea is simply based on the observation that the performance of class i along its route \mathcal{R}_i in a tree can be approximated by the performance of class 1 in Figure 1, where the links $\{C_1, \dots, C_n\}$ correspond to the links in \mathcal{R}_i . Applying the parking lot configuration all the way to the access link of an individual user in class i implies that the access link is treated as a link with rate a_i and the load on the access link equals zero. This is because our underlying assumption in the model is that arrivals are Poisson implying an infinite user population and hence the load of an individual user is zero.

More precisely, the application of the parking lot configuration to approximate the throughput of class i flows in a tree is the following. Assume that the links constituting the route for class i , \mathcal{R}_i , are arranged in a list in the order as they exist in the tree starting from the leaf and ending with the root link. We use the notation $\mathcal{R}_i(m)$ to denote the m th element in \mathcal{R}_i . Leaf link here refers to the first link that is shared by the users of a given class i . The access link a_i of an individual user corresponds to link C_1 in Figure 1, and the load on that link is zero. The link C_2 in Figure 1 corresponds to the leaf link of class- i flows, i.e., the link $\mathcal{R}_i(1)$, and the load on that link equals $R_{\mathcal{R}_i(1)}$, etc. Then the parking lot approximation for

the throughput of class- i flows, γ_i^{PL} can be expressed as

$$\gamma_i^{\text{PL}} = \left(\frac{1}{a_i} + \left(\frac{1}{C_{\mathcal{R}_i(1)} - \rho_i} - \frac{1}{C_{\mathcal{R}_i(1)}} \right) + \sum_{m=2}^{|\mathcal{R}_i|} \left(\frac{1}{C_{\mathcal{R}_i(m)} - R_{\mathcal{R}_i(m)}} - \frac{1}{C_{\mathcal{R}_i(m)} - R_{\mathcal{R}_i(m-1)}} \right) \right)^{-1}, \quad (7)$$

4 Network dimensioning using balanced fairness

We propose using models based on the notion of balanced fairness for network dimensioning. Among the appealing properties of these models are the insensitivity property (robustness) and the fact that performance can be numerically solved efficiently. From the point of view of dimensioning, another useful property is that the performance measure captured by these models is a fundamental end-to-end metric that is also meaningful for any application, namely the average throughput that flows observe along a given route.

4.1 Problem decomposition

The dimensioning problem is first decomposed into two parts: (i) dimensioning of access networks (ANs), and (ii) dimensioning of core networks to which the access networks are connected, see Figure 2. The idea is that the dimensioning of these different parts of the network can be done in a decoupled manner.

In the core networks, such as the ISP core network in Figure 2, the capacities of the links are much greater than the capacities of the access links of users due to the significant amount of traffic aggregation. When the offered load to the (core) links exceeds the access rate of any user by a large factor (an order of magnitude, say) it holds generally that a moderate over-provisioning of the core links is sufficient to render the core network transparent. In the case of balanced fairness, this can be easily seen in the properties of the store-and-forward

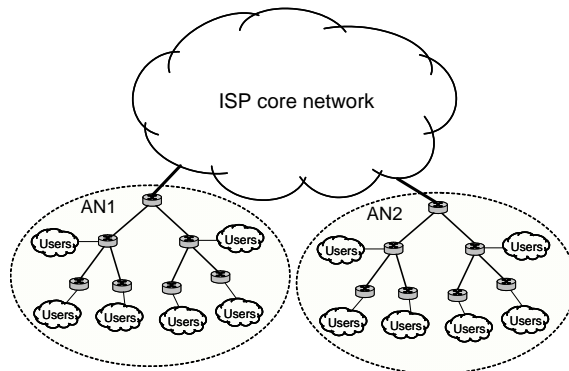


Figure 2: Decomposition of the network dimensioning problem.

bound (5). The delay contributions from the links along the route are inversely proportional to the amount of available bandwidth, i.e., the terms $(C_l - R_l)$. Even under moderate over-provisioning (by a factor of 20%, say) the amount of available bandwidth in the core links makes the delay contributions negligible compared to the delays caused by the links closer to a particular user's access, where the available bandwidth is in magnitude closer to the access rate limitation. In the core, moderate over-provisioning should be mostly easy as fiber is already installed in the core networks and capacity upgrades simply imply that interface cards are changed in the core routers.

On the other hand, fiber has not typically been installed in access networks, and thus capacity there may be scarce. Changing the existing copper wire to fiber cannot be done that easily, either. Therefore, in the access networks a more detailed dimensioning based on a meaningful end-to-end performance measure, such as mean flow throughput, should be performed, and to this end we propose using models based on balanced fairness.

4.2 Optimal dimensioning of tree-type access networks

To have a unique solution for the capacities C_l , the problem needs to be formulated as an optimization problem. We assume that each class in the network is given a minimum criterion for the flow throughput denoted by γ_i^{\min} . Additionally, the load of each class ρ_i is known, as well as the topology of the tree. The classes may have access rate limits a_i or not, depending on the considered scenario. Associated with each link is a constant b_l representing the cost per bit for link l . The idea is to only optimize with respect to the network links C_l and the access rate limits are not subject to optimization but rather they are considered as given parameters, i.e., as properties of a given user population (users are grouped according to their access rate).

The optimization problem for determining the link capacities such that the given throughput criteria are met can be defined as

$$\begin{aligned} \min \quad & \sum_{l \in \mathcal{L}} b_l C_l \\ \text{subject to} \quad & \sum_{i \in \mathcal{F}_l} \rho_i < C_l, \quad \forall l \in \mathcal{L}, \\ & \gamma_i \geq \gamma_i^{\min}, \quad \forall i = 1, \dots, N. \end{aligned} \tag{8}$$

In (8) the first set of constraint equations correspond to the stability requirements and the second set of constraint equations define the region of feasible solutions. Essentially, we are requiring that at the optimal solution all classes achieve at least a given minimum throughput.

Alternatively, the throughput constraints in (8) can be also expressed with respect to the mean flow-level transfer delays by recalling that $\gamma_i = E[S_i]/E[T_i]$. Thus, the optimization

problem can be expressed in an equivalent form

$$\begin{aligned} \min \quad & \sum_{l \in \mathcal{L}} b_l C_l \\ \text{subject to} \quad & \sum_{i \in \mathcal{F}_l} \rho_i < C_l, \quad \forall l \in \mathcal{L}, \\ & \mathbb{E}[T_i] \leq \mathbb{E}[S_i]/\gamma_i^{\min}, \quad \forall i = 1, \dots, N. \end{aligned} \tag{9}$$

In fact, since $\mathbb{E}[T_i] = \mathbb{E}[S_i]/\gamma_i$ the inequality constraint $\mathbb{E}[T_i] \leq \mathbb{E}[S_i]/\gamma_i^{\min}$ is the same as $1/\gamma_i \leq 1/\gamma_i^{\min}$.

The constraints in (8) and (9) are convex, and when using the parking lot approximation or the store-and-forward bound the constraints are also continuously differentiable. However, the exact solution to the throughput is not everywhere continuously differentiable. Also, in numerical studies the use of (8) or (9) depends on whether it is more convenient to work with throughput or the mean flow transfer delays. When computing the throughput exactly from the normalization constant G , the throughput is obtained from 4 and in that case the optimization is carried out by solving (8). However, if we employ the store-and-forward bound for the throughput, then the optimization is numerically more efficient when using the formulation (9). Finally, note that the objective cost function can easily be replaced with some other than a linear one, for example the cost could be proportional to \sqrt{C} modeling the fact that cost per bit increases slower for larger capacities.

5 Numerical examples

In the numerical examples we will use the store-and-forward bound and/or the parking lot approximation to approximate the throughput (flow transfer delay). This makes the solution of the optimization problems numerically very efficient, and by using the exact algorithms we can compare how much over-provisioning the use of the approximations causes.

5.1 Example with a simple network

First we experiment with a simple n -branch tree network, as illustrated in Figure 3. The network consists of a root link C_0 and n branches each with a capacity C_1 , and there are no per-flow access rate limitations. The load in each class is the same $\rho_i = \rho, \forall i = 1, \dots, n$, and also the minimum throughput criterion $\gamma_i^{\min} = \gamma^{\min}, \forall i = 1, \dots, n$.

Using the store-and-forward bound in optimization problem (9), allows us to solve the problem explicitly. The Lagrangian function is given by

$$L = b_0 C_0 + n b_1 C_1 + \beta \left(\frac{1}{C_0 - n\rho} + \frac{1}{C_1 - \rho} - \frac{1}{\gamma^{\min}} \right),$$

where β is the Lagrangian multiplier associated with the constraint equation. At the optimum the inequality constraint is satisfied as an equality, since costs are minimal when the

performance criterion is satisfied as an equality. The optimal solution is then found to be

$$\begin{cases} C_0 &= n\rho + \gamma^{\min} \left(1 + \sqrt{n} \sqrt{\frac{b_1}{b_0}}\right), \\ C_1 &= \rho + \gamma^{\min} \left(1 + \frac{1}{\sqrt{n}} \sqrt{\frac{b_0}{b_1}}\right). \end{cases}$$

Next we discuss the above for $b_0 = b_1 = 1$. It is easy to see that asymptotically, as $n \rightarrow \infty$, the optimal dimensioning is given by $C_0 = n\rho + \gamma^{\min}$ and $C_1 = \rho + \gamma^{\min}$, and the optimal total cost is then $C_0 + nC_1 = 2n\rho + (n+1)\gamma^{\min}$. This means that when the number of leaf links n is large, the bottleneck C_0 is dimensioned such that the link capacity equals the bottleneck load $n\rho$ and a little bit extra (γ^{\min}), and for the access links it is sufficient to have a bandwidth where the “remaining” capacity $C_1 - \rho$ equals the dimensioning criterion γ^{\min} . The rate at which the asymptotical solutions are approached are illustrated in Figure 4. The left and center figures depict C_0/n (left figure) and C_1 (middle figure), respectively, as a function of n for $\rho = 1$ Mbit/s and for various values of the target dimensioning criterion, $\gamma^{\min} = \{0.5, 5, 10\}$ Mbit/s (from bottom up). Additionally, the right figure shows the total cost of the network $(C_0 + nC_1)/n$ as a function of n . The dimensioning criterion $\gamma^{\min} = 10$ Mbit/s means that the throughput of a flow is much higher than the offered load (bandwidth of the leaf is not shared much among competing flows, resembling an access link of a single user). Conversely, when $\gamma^{\min} = 0.5$ Mbit/s the throughput of a flow is close to the offered load (bandwidth sharing occurs also at the leaf link). As can be seen, for C_1 the asymptotical region is approached quickly and already with $n = 20$ the solution is close to the asymptotical values. However, for C_0 the convergence is slower. Note that the shape of the curves does not depend on ρ , i.e., the value of ρ just scales the curves up or down.

Finally, we examine the amount of over-provisioning caused by use of the conservative store-and-forward bound for evaluating the performance, instead of computing it exactly. To this end, we compare the performance as given by the exact solution when using the link bandwidths resulting from the optimization (with the use of the store-and-forward bound) to the target dimensioning criterion (which is fulfilled exactly in the optimization). The resulting over-provisioning factor, $(\gamma(\mathbf{C}^*) - \gamma^{\min})/\gamma^{\min}$, where $\gamma(\mathbf{C}^*)$ denotes the exact solution with link bandwidths from a given instance of the optimization problem, is illustrated in Figure 5 for $\rho = 1$ Mbit/s as a function of n for various γ^{\min} . As can be seen the error is greater for smaller values of n and smaller values of γ^{\min} (higher degree of bandwidth sharing). Again note that the curves only depend on the ratio γ^{\min} .




Figure 3: Homogeneous n -branch tree.




Figure 4: Scaled required capacities for C_0 (left) and C_1 (center), and the total cost (right) as a function of n for various target dimensioning criteria.




Figure 5: Amount of over-provisioning caused by using the store-and-forward bound in the dimensioning.

5.2 Example with a larger network and access rate limits

Here we consider dimensioning of a larger multi-level tree where the individual flows are also subject to a per-flow access rate limitation corresponding to current ADSL access speeds. The topology of the tree is illustrated in Figure 6. The network consists of seven links and there are four access routers. To each access router two types of customers have been connected, with access speeds equal to 256 kbits/s and 512 kbit/s. Altogether there are eight traffic classes such that $a_{i,1} = 256$ kbit/s for $i = 1, \dots, 4$, and $a_{i,2} = 512$ kbit/s for $i = 1, \dots, 4$. To simplify the situation, it is assumed that the traffic load is homogenous, i.e., $\rho_{i,j} = \rho, \forall i = 1, \dots, 4$ and $j = 1, 2$. Note that in this case the performance for classes (1, 1) and (2, 1) is identical, and similarly for classes (1, 2) and (2, 2).

Some numerical results on dimensioning of the above network are illustrated in Figure 7 when using the store-and-forward bound. In the results, we have chosen the target dimensioning criterion to be a given fraction α of the access rate, i.e., $\gamma_{i,j}^{\min} = \alpha \cdot a_{i,j}$. The figure shows the optimal dimensioning for link C_1 (left), C_7 (middle) and the total cost of the network (right) as a function of the fraction parameter $\alpha \in [0.1, 0.9]$ for $\rho = 1, 2, 3$ Mbit/s.

The same results as above but when using the parking lot approximation are given in Figure 8. From the figure we can see that the general level of the curves is lower than in Figure 7, where the store-and-forward bound has been used. This indicates that with the parking lot approximation the dimensioning more closely matches the results as would be obtained using the exact solution for the performance.




Figure 6: General tree example.




Figure 7: Optimal capacity for C_1 (left), C_7 (middle) and the total cost of the network (left) when using the store-and-forward bound.

Note that for the considered network the solutions to the optimization problem are determined by the constraints related to the classes with the higher access rate limits, i.e., classes (1, 2), (2, 2), (3, 2) and (4, 2). Thus, only constraints pertaining to these are satisfied at the optimum as strict equalities. The over-provisioning resulting from the use of the conservative store-and-forward bound compared to the parking lot approximation are illustrated in Figure 9 as a function of the target throughput criterion (as controlled by the fraction parameter). Left figure shows the results for the classes (1, 2) and (4, 2). The results for the parking lot approximation are shown with solid lines and the store-and-forward bound is shown with dashed lines. The lowest solid diagonal line represents the target dimensioning criterion, i.e., the result if the dimensioning criterion and the actual performance are equal. As can be seen, the parking lot approximation indeed results in a tighter dimensioning than with the store-and-forward bound. For completeness, the right figure shows the same results for classes (1, 1) and (4, 1) for which the constraints in the optimization are not satisfied as equalities. Hence, it can be seen that the over-provisioning margin is even greater for these classes. However, still the parking lot approximation yields a tighter dimensioning result.




Figure 8: Optimal capacity for C_1 (left), C_7 (middle) and the total cost of the network (left) when using the parking lot approximation.




Figure 9: Comparison of the over-provisioning caused by the use of store-and-forward bound and the parking lot approximation for classes with $a_{i,j} = 512$ kbit/s (left) and $a_{i,j} = 256$ kbit/s (right).

6 Conclusions

Access networks are the part of the network where users directly connect to the network, and for these networks it is important to base the dimensioning on some, from the user's point of view, meaningful QoS criterion. In this paper, we have proposed the use of models based on the notion of balanced fairness for dimensioning access networks that carry elastic traffic (e.g., access networks to the public Internet). The notion of balanced fairness provides an elegant abstraction for the way that stochastically arriving and departing flows share the bandwidth of the network. It serves as a computationally tractable model of how practical bandwidth sharing algorithms, such as TCP, perform. Models based on balanced fairness are also robust in the sense that the performance depends only on the load of the system. Furthermore, the dimensioning can be based on a natural flow-level QoS requirement, i.e., the per-flow throughput. Using the store-and-forward bound for approximating the performance, the optimization problem is also numerically simple yielding a conservative estimate of the required capacity. Even tighter estimates can be obtained by using the parking lot approximation. Efficient exact solution algorithms can be used to verify the amount of over-provisioning resulting from the use of the store-and-forward bound or the parking lot approximation. The numerical examples demonstrate the viability of the proposed approach.

Open issues for future research include the following. In access networks, the assumption of an infinite user population is not always well justified. The models based on balanced

fairness can be defined with a finite user population but computation of the performance becomes more complex. Approximations on a single link with a finite user population have been given in [2], but similar approximations would be useful for tree networks, as well. Additionally, modern networks include both streaming and elastic traffic, and a full account of the dimensioning problem should consider both types of traffic.

References

- [1] D. Bertsekas and R. Gallager. *Data networks, 2nd edition*. Prentice-Hall, second edition, 1992.
- [2] T. Bonald, P. Olivier, and J. Roberts. Dimensioning high speed IP access networks. In *Proceedings of 18th International Teletraffic Congress*, pages 241–251, August 2003.
- [3] T. Bonald and A. Proutière. Insensitive bandwidth sharing in data networks. *Queueing Systems*, 44:69–100, 2003.
- [4] T. Bonald and A. Proutière. On performance bounds for balanced fairness. *Performance Evaluation*, 55:25–50, 2004.
- [5] T. Bonald, A. Proutière, J. Roberts, and J. Virtamo. Computational aspects of balanced fairness. In *Proceedings of 18th International Teletraffic Congress*, pages 801–810, August 2003.
- [6] T. Bonald and J. Virtamo. Calculating the flow level performance of balanced fairness in tree networks. *Performance Evaluation*, 58:1–14, 2004.
- [7] J. Charzinski. Fun factor dimensioning for elastic traffic. In *Proceedings of ITC Specialist Seminar on Internet Traffic Measurement, Modeling and Management*, September 2000.
- [8] J.W. Cohen. The multitype phase service network with generalized processor sharing. *Acta Informatica*, 12:245–284, 1979.
- [9] G. Davies, M. Hardt, and F. Kelly. Network dimensioning, service costing and pricing in a packet switched environment. *Telecommunications Policy*, 28:391–412, 2004.
- [10] A. K. Erlang. Solution of some problems in the theory of probabilities of significance in automatic telephone exchanges. *Elektroteknikerer*, 13, 1917.
- [11] F. Kelly, A. Maulloo, and D. Tan. Rate control in communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research Society*, 49:237–252, 1998.
- [12] G. Malicskó, G. Fodor, and M. Pióro. Link capacity dimensioning and path optimization for networks supporting elastic services. In *Proceedings of IEEE International Conference on Communications (ICC)*, pages 2304 – 2311, April 2002.

- [13] P. Nilsson and M. Pióro. Fair bandwidth allocation and links' dimensioning in elastic traffic networks. In *Proceedings of Nordic Teletraffic Seminar (NTS-16)*, pages 56 – 69, August 2002.
- [14] P. Nilsson and M. Pióro. Solving dimensioning tasks for proportionally fair networks carrying elastic traffic. *Performance Evaluation*, 49:371–386, 2002.
- [15] M. Pióro and D. Medhi. *Routing, Flow, and Capacity Design in Communication and Computer Networks*. Morgan Kaufmann Publishers, 2004.
- [16] T. Pioró, G. Malicskó, and G. Fodor. Optimal link capacity dimensioning in proportionally fair networks. In *Proceedings of Networking 2002*, pages 277–288, May 2002.
- [17] A. Riedl, M. Perske, T. Bauschert, and A. Probst. Dimensioning of ip access networks with elastic traffic. In *Proceedings of 9th International Telecommunication Network Planning Symposium (Networks 2000)*, September 2000.
- [18] Vesa Timonen. Simulation studies on performance of balanced fairness. Master's thesis, Helsinki University of Technology, 2003.
- [19] H. van den Berg, M. Mandjes, R. van de Meent, A. Pras, F. Roijers, and P. Venemans. QoS-aware bandwidth provisioning for IP network links. submitted for publication, February 2004.