

Dimensioning methods for data networks with flow-level QoS requirements

Pasi Lassila and Aleks Penttinen and Jorma Virtamo
Helsinki University of Technology
P.O.Box 3000, FIN-02015 TKK Finland
Email: {Pasi.Lassila, Aleks.Penttinen, Jorma.Virtamo}@tkk.fi

Abstract—We consider dimensioning of data networks. The network is modeled in a dynamic setting where elastic flows (file transfers) arrive randomly and share the bandwidth according to balanced fairness. Simple methods are derived for determining the link capacities so that given flow-level throughput requirements are satisfied. We consider two different approaches to define the throughput requirements: single constraint on average throughput in the network and separate constraint for every route. The results enable a simple characterization of the order-of-magnitude of the required capacities, which can be utilized in practical network planning and dimensioning.

I. INTRODUCTION

Optimal network dimensioning is a classical problem in teletraffic theory. The role of dimensioning is to assist network planning by providing reasonable estimates of the required resources to fulfill a certain demand with a predefined level of service. For circuit switched networks the dimensioning criterion has traditionally been the call blocking probability, and the related dimensioning methods have been based on Erlang's formula [1]. One attractive property of Erlang's formula is its insensitivity property, i.e., the performance of the system depends only on the traffic load and not on any detailed properties of call holding times. This makes the application of the formula for dimensioning simple, since only a single parameter is needed to characterize the traffic.

Here our focus is on dimensioning of wireline data networks, such as IP networks, which carry mostly data traffic consisting of file transfers. Data traffic is elastic by nature, meaning that the applications transmitting the data can sustain fluctuations in the transmission rate (cf., operation of TCP). For data traffic it is not the delays of individual packets that reflect the QoS experienced by a user. Instead, the user-experienced performance manifests itself at the so called *flow level*.

At the flow level, we have a dynamic system, where flows (i.e., file transfers) arrive randomly and depart upon completion. During a file transfer the flow shares the bandwidth with other flows in the system and, accordingly, the rates of the flows vary randomly over time. The performance at the flow-level is expressed as the mean transmission rate during a file transfer, i.e., the *per-flow throughput*, or mean flow transfer delay (or roughly mean file transfer delay). The dimensioning problem for fixed data networks is to determine the link capacities of the network for a given traffic pattern so

that the per-flow throughput experienced by users remains at acceptable levels.

To elucidate, consider a single link with an average traffic load, say, 8 Mbps. The dimensioning problem could be to determine the capacity of the link so that the average data rate of file transfers using the link would be 1 Mbps. By modeling the link using the well-known M/G/1-PS queue we easily get the result that a link with capacity 9 Mbps suffices. However, the situation becomes a lot trickier in a network setting where the flows traverse through several links with different loads. This is the topic of this paper.

The most important aspect affecting the flow-level performance is the bandwidth sharing policy. Common idealized bandwidth sharing schemes are max-min fairness [2], proportional fairness [3], and more recently, balanced fairness [4]. In this work, we apply models based on balanced fairness due to the following reasons. Under balanced fairness computing the throughput is significantly easier than with other sharing schemes in the dynamic setting. Furthermore, the throughput depends only on traffic load (insensitivity), i.e., similarly as in Erlang's formula, knowledge of only one traffic parameter is needed when applying the results for dimensioning. The other sharing models require more detailed traffic modeling. However, balanced fairness can also be used to approximate both max-min and proportional fairness [5]. Even under balanced fairness throughput evaluation becomes computationally difficult in networks of realistic size, but for dimensioning purposes suitable bounds are available [6], [7].

To define the dimensioning problem we assume that the resources are shared according to balanced fairness, the network topology is given, routes are fixed and the offered load together with the throughput requirements are known. To start with, the link capacities need to be such that the network is able to carry the traffic, i.e., that the corresponding dynamic system is stable. At this stability limit the network is not able to support any flow-level performance. Thus, our focus is on how much *excess* capacity is needed to satisfy given throughput constraints.

We consider two different approaches to define the throughput constraints. In the first, we set a constraint on the average throughput in the network. In the second, we constrain the throughput of each class (i.e., route) separately. Our contributions are:

- a) In the average throughput setting, we provide an explicit

formula for an upper bound on the excess capacity of the links.

- b) In the per-class throughput setting, we provide an explicit formula for a lower bound on the excess capacity of the links. For the upper bound, an efficient iterative solution method is presented.

The results enable a simple characterization of the order-of-magnitude of the required capacities, which can be utilized in practical network planning and dimensioning.

The paper is organized as follows. In Section II we review the related literature. Section III outlines the network model. Section IV defines the dimensioning problems. Per-class throughput problems are addressed in Section V. We illustrate the application of the methods via some numerical examples in Section VI. Finally, the conclusions are given in Section VII.

II. RELATED WORK

Existing methods in the literature on dimensioning of 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. In static models, the dimensioning problem is often formulated as a multi-commodity optimization problem, see [8]–[10]. The network consists of routes 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 according to various fairness criteria (max-min or proportional fairness). Additionally, other considerations may easily be included, such as network failures. In principle, the multi-commodity flow problem is well defined and easy to apply for various types of network design problems. However, the approach does neglect one important aspect, namely the dynamic and stochastic nature of the traffic.

A classic data network dimensioning method that is also formulated as an optimization problem, but employing a stochastic traffic model, is the square-root method in [2]. 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 the approach is that the performance criterion is a packet-level metric which is not a meaningful end-to-end performance measure for elastic data traffic.

There are also a number of papers that focus on dimensioning of a single bottleneck link. In [11], different models (on-off fluid model, processor sharing and Brownian motion) have been used to evaluate the dimensioning requirements given by different single link traffic models. The results are also compared against measurements. The so called GPS (generalized processor sharing) model [12] has been applied for dimensioning of elastic TCP traffic in [13]. The GPS model

assumes an infinite user population (Poisson arrivals) and 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 equally. A finite user population variant of GPS has been studied in [14], where a simple and explicit formula is derived for the required capacity, given the number of users, offered traffic and a target dimensioning criterion for the mean useful rate of a user.

In our previous work [15] we considered the dimensioning of wireless mesh networks under flow-level QoS requirements. The problem addressed in the present work appears as a sub-problem in an approximation scheme for mesh network dimensioning, although the two problems are different.

Similarly as in [14] and [13], we apply flow-level models, but instead of focusing on a single link, we consider the network case and apply models based on balanced fairness for dimensioning of data networks. Our formulation is in spirit similar to the one in [2], except that the performance measure in our case is not the packet delay, but the flow level throughput associated with file transfers in a network.

III. FLOW-LEVEL MODELING OF ELASTIC TRAFFIC

Consider a network of L links with traffic routed along N routes. Each route corresponds to a flow class. The routing is described by the matrix \mathbf{A} , where the element $a_{il} = 1$ if class i uses link l and 0 otherwise. Let \mathcal{R}_i be the set of links used by class i . When link l is on route \mathcal{R}_i we use the notation $l \in \mathcal{R}_i$. Correspondingly, the notation \mathcal{F}_l is used for the set of flow classes using link l .

The flows arrive randomly and have finite durations in each flow class. The traffic of class i is characterized by the traffic load ρ_i , which is the product of the mean arrival rate of class- i flows, λ_i , and their mean flow size, $E[S_i]$. The vector of loads is denoted by $\boldsymbol{\rho}$. The link loads are given by the vector $\mathbf{r} = \mathbf{A}^T \boldsymbol{\rho}$. The l th component of \mathbf{r} is denoted by r_l . All vectors are assumed to be column vectors.

We assume that the bandwidth is shared dynamically among the on-going flows and that the sharing policy is balanced fairness [4]. Balanced fairness is a resource sharing notion that essentially renders the flow-level traffic process reversible. The stationary distribution of the system depends only on the traffic loads under rather general assumptions. This significantly facilitates analysis of the system, while the performance remains comparable with other fair sharing schemes [5]. Balanced fairness has been extensively applied in performance analysis of various communication networks, see, e.g., [16]–[19].

The performance measure we consider is the flow throughput of class i , which is defined as the ratio of mean flow size and flow duration in the class, i.e.,

$$\gamma_i = \frac{E[S_i]}{E[T_i]}.$$

Under balanced fairness one can obtain explicit formulas for γ_i for particular systems [4], [16], [20], but generally the approach results in a recursive formulation which is too

cumbersome for dimensioning purposes. For dimensioning we can utilize the upper and lower bounds for throughput, introduced in [6] and [7].

The most straightforward approach to the throughput of class- i flows is to approximate the throughput by the available capacity of the bottleneck link along the route of the class. In this case, the throughput function $\Gamma_i(\mathbf{c})$ is given by

$$\Gamma_i(\mathbf{c}) = \left(\max_{l \in \mathcal{R}_i} \frac{1}{c_l - r_l} \right)^{-1}, \quad (1)$$

where \mathbf{c} denotes the vector of link capacities and c_i its i th component. In fact, (1) provides an upper bound on the throughput, as shown in [6], and it is accurate when the performance of the class is mostly determined by a single bottleneck.

As a lower bound for the throughput, the so called store-and-forward (SF) bound has been proposed in [6]. The bound is based on assuming that the flow is transmitted along the route so that the flow is first received entirely by the next hop node and only then forwarded to the next node along the route (hence the name store-and-forward). The throughput function $\Gamma_i(\mathbf{c})$ in this case equals

$$\Gamma_i(\mathbf{c}) = \left(\sum_{l \in \mathcal{R}_i} \frac{1}{c_l - r_l} \right)^{-1}. \quad (2)$$

Recently, however, Bonald improved this bound in [7]. The corresponding throughput function $\Gamma_i(\mathbf{c})$ is

$$\Gamma_i(\mathbf{c}) = \left(\max_{l \in \mathcal{R}_i} \frac{1}{c_l} + \sum_{l \in \mathcal{R}_i} \frac{r_l}{c_l} \frac{1}{c_l - r_l} \right)^{-1}. \quad (3)$$

We refer to (3) as the improved store-and-forward (ISF) bound. ISF is a considerably tighter bound in certain cases than SF, but dimensioning using the SF bound is computationally easier, as will be discussed later.

IV. THE DIMENSIONING PROBLEM FORMULATIONS

The problem of dimensioning is to determine the required amount of bandwidth on each link to fulfill a given throughput requirement when the routing is fixed by the matrix \mathbf{A} and the load vector $\boldsymbol{\rho}$ is given. Note that in the dimensioning $\boldsymbol{\rho}$ represents an estimate of the traffic demand in the network.

The stability of the dynamic system requires that the link capacities \mathbf{c} are greater than \mathbf{r} . At this limit the flow-level performance approaches zero. Thus, more capacity is needed to attain any given flow-level performance objectives. We simplify the notation here by considering only this excess capacity $\mathbf{d} = \mathbf{c} - \mathbf{r}$.

A. Dimensioning based on average throughput

When considering how much extra capacity is required, it is natural to formulate the dimensioning as an optimization problem, where the objective is to minimize the overall cost of the network subject to given performance requirements. The performance requirement can be chosen as the average

throughput of the flows in the network, $\gamma_{\text{ave}} = E[S]/E[T]$, where $E[S]$ is the mean flow size,

$$E[S] = \frac{1}{\sum_i \lambda_i} \sum_i \lambda_i E[S_i],$$

and $E[T]$ is the mean flow transfer delay in the network. By using the store-and-forward bound, an upper bound for the mean flow transfer delay in the network is given by the function

$$E[T] = \sum_i \frac{\lambda_i}{\sum_j \lambda_j} \sum_{l \in \mathcal{R}_i} \frac{E[S_i]}{d_l},$$

Thus, we have the throughput function

$$\Gamma_{\text{ave}}(\mathbf{d}) = \frac{E[S]}{E[T]} = \left(\sum_i \frac{\rho_i}{\rho_{\text{tot}}} \sum_{l \in \mathcal{R}_i} \frac{1}{d_l} \right)^{-1},$$

where $\rho_{\text{tot}} = \sum_i \rho_i$. The associated optimal network dimensioning problem can be expressed as

$$\begin{aligned} \min_{\mathbf{d}} \mathbf{w}^T \mathbf{d}, \\ \Gamma_{\text{ave}}(\mathbf{d}) \geq \gamma_{\text{ave}}, \\ \mathbf{d} > 0, \end{aligned} \quad (4)$$

where \mathbf{w} is the vector of cost per unit capacity on each link and γ_{ave} is our target average per-flow throughput in the network. The cost vector \mathbf{w} is included for generality.

This optimization problem can be solved explicitly. The Lagrangian function Θ associated with (4) is given by

$$\Theta = \sum_l w_l d_l + \lambda \left(\sum_i \frac{\rho_i}{\rho_{\text{tot}}} \sum_{l \in \mathcal{R}_i} \frac{1}{d_l} - \frac{1}{\gamma_{\text{ave}}} \right), \quad (5)$$

where λ is the Lagrangian multiplier of the average throughput constraint. The derivative of (5) goes to zero at

$$d_l = \sqrt{\frac{\lambda r_l}{w_l \rho_{\text{tot}}}}, \quad \forall l.$$

At the optimum the throughput constraint is satisfied as an equality, which yields

$$\sqrt{\lambda} = \gamma_{\text{ave}} \sum_l \sqrt{\frac{r_l w_l}{\rho_{\text{tot}}}}.$$

Thus, the optimal values of d_l are given by

$$d_l = \gamma_{\text{ave}} \sqrt{\frac{r_l}{w_l \rho_{\text{tot}}}} \sum_j \sqrt{\frac{r_j w_j}{\rho_{\text{tot}}}}, \quad \forall l. \quad (6)$$

The result above can be seen as the flow-level counterpart of the original square-root method [2] for dimensioning according to the packet-level delay. Note that such a simple explicit solution for this formulation is available only when applying the store-and-forward bound.

B. Dimensioning based on per-class throughput

Due to the explicit solution, the above method is indeed easy to apply in practise for any network and it can be used to provide quickly computable estimates of the required capacities. However, the solution does not provide any guarantees on the per-class performance in the system, i.e., in the obtained solution some classes may have significantly lower throughput than others. This motivates us to consider also optimization problems, where the throughput requirement is defined separately for each class.

The dimensioning problem is then to fix the link capacities so that the throughput requirements on each route are satisfied and that the overall cost is minimized. Formally, we search for \mathbf{d} which solves the following problem:

$$\begin{aligned} \min_{\mathbf{d}} \mathbf{w}^T \mathbf{d}, \\ \Gamma_i(\mathbf{d}) \geq \gamma_i, \quad \forall i, \\ \mathbf{d} > 0, \end{aligned} \quad (7)$$

where γ_i denotes the per-class throughput requirement and \mathbf{w} is a positive capacity cost vector. In (7), the throughput function $\Gamma_i(\mathbf{d})$ can be any of the throughput functions (1-3).

By applying (1), a simple lower bound on the required excess capacities (and the network cost) can be obtained. This approximation gives us the following constraints on the link capacities

$$\left(\max_{l \in \mathcal{R}_i} \frac{1}{d_l} \right)^{-1} \geq \gamma_i, \quad \forall i.$$

Since the constraint holds for the maximum value in the class, it must hold for all the links in the class, i.e.,

$$\frac{1}{d_l} \leq \gamma_i^{-1}, \quad \forall l \in \mathcal{R}_i, \forall i.$$

These constraints are satisfied if the most stringent ones for each d_l are satisfied and we get

$$d_l = \max_{i \in \mathcal{F}_l} \gamma_i, \quad \forall l. \quad (8)$$

Thus, we have an explicit solution for the required link capacities, $c_l = r_l + d_l$, where d_l represents an optimistic estimate (lower bound) of the excess capacity needed on link l to satisfy the throughput requirements of all classes. Note that this result can be interpreted so that the network is dimensioned assuming that each link l behaves as an independent M/G/1-PS queue with an offered load r_l .

By applying the throughput functions (2) or (3) upper bounds on the required capacities can be obtained. In the next section, we propose efficient methods to solve (7) when throughput function (2) or (3) is used.

V. UPPER BOUND ON LINK CAPACITIES WITH PER-CLASS REQUIREMENTS

The dimensioning problem (7) with throughput formulas (2) and (3) is a non-linear optimization problem that can be tedious to solve, especially for large problem instances. One of our main contributions is the following scalable iterative method for solving the optimization problem.

A. Store-and-forward bound

The store-and-forward (SF) bound gives us the throughput function

$$\Gamma_i(\mathbf{d}) = \left(\sum_{l \in \mathcal{R}_i} d_l^{-1} \right)^{-1}. \quad (9)$$

Thus, dimensioning problem (7) for the excess capacity \mathbf{d} becomes:

$$\begin{aligned} \min_{\mathbf{d}} \mathbf{w}^T \mathbf{d} \\ \mathbf{A} \mathbf{d}^{-1} \leq \boldsymbol{\gamma}^{-1}, \\ \mathbf{d} > 0. \end{aligned} \quad (10)$$

where we use the notation $\mathbf{d}^{-1} = (d_1^{-1}, \dots, d_L^{-1})^T$ (also for $\boldsymbol{\gamma}$) for brevity. The Lagrangian dual to this problem is given by

$$\begin{aligned} \max_{\mathbf{u}} \Theta(\mathbf{u}) \\ \mathbf{u} \geq \mathbf{0}, \end{aligned} \quad (11)$$

where

$$\Theta(\mathbf{u}) = \inf_{\mathbf{d} > 0} \{ \mathbf{w}^T \mathbf{d} + \mathbf{u}^T (\mathbf{A} \mathbf{d}^{-1} - \boldsymbol{\gamma}^{-1}) \}. \quad (12)$$

Although the dual problem is tractable by standard numerical optimization techniques [21], the special structure of the problem allows a simple iterative solution. Note first that the vector \mathbf{d} realizing the infimum in (12) is given by

$$\mathbf{d} = \sqrt{\frac{\mathbf{A}^T \mathbf{u}}{\mathbf{w}}}, \quad (13)$$

where the division by \mathbf{w} stands for component-wise division. Recall that in the optimum the complementary slackness (CS) condition,

$$u_i \left(\sum_{l \in \mathcal{R}_i} d_l^{-1} - \gamma_i^{-1} \right) = 0,$$

must hold for all classes. By utilizing the relation (13) the CS conditions result in a natural iteration for \mathbf{u} . Start from any strictly positive \mathbf{u}_0 and iterate the value until convergence using

$$\mathbf{u}_k = \mathbf{u}_{k-1} * \boldsymbol{\gamma} * \mathbf{A} \sqrt{\frac{\mathbf{w}}{\mathbf{A}^T \mathbf{u}_{k-1}}}. \quad (14)$$

The symbol $*$ is used to denote component-wise multiplication.

Extensive numerical experiments suggest that the iteration converges. Note that if the iteration converges, the limiting value is indeed the optimal solution to (11). Given that the iteration is started from a strictly positive \mathbf{u}_0 , the form of (14) ensures that $\mathbf{u}_k > 0$ for all k , i.e., all \mathbf{u}_k are feasible solutions to the dual. Assume that the iteration converges to a limit value \mathbf{u}_* . It is easy to see that then the limiting coefficient of \mathbf{u}_{k-1} in (14), $\boldsymbol{\gamma} * \mathbf{A} \mathbf{d}^{-1}$, cannot be larger than 1 for any class. This is equivalent to that the corresponding limiting value of (13), \mathbf{d}_* , is a feasible solution to (10). Finally, at the limit the CS conditions are satisfied. Thus, \mathbf{d}_* and \mathbf{u}_* are optimal solutions to (10) and (11), respectively.

In conclusion, the solution to the dimensioning problem is given by (13) where \mathbf{u} is obtained from the iteration (14).

B. Improved store-and-forward bound

The throughput function of the improved store-and-forward bound (ISF) is given by

$$\Gamma_i(\mathbf{d}) = \left(\max_{l \in \mathcal{R}_i} \frac{1}{d_l + r_l} + \sum_{l \in \mathcal{R}_i} \frac{r_l}{d_l(d_l + r_l)} \right)^{-1}. \quad (15)$$

With a suitable restatement of the constraints the solution process is analogous to the store-and-forward case. To eliminate the max-operations from the constraints we form two matrices \mathbf{B}_1 and \mathbf{B}_2 as follows. \mathbf{B}_1 is generated from \mathbf{A} by duplicating each row i of \mathbf{A} as many times as there are non-zero elements on that row (i.e., the number of links used by route i). \mathbf{B}_2 is generated from \mathbf{A} by taking each of the non-zero elements to its own row (which is otherwise full of zeros). As with \mathbf{B}_1 , the right hand side of the constraints are obtained by duplicating the elements γ_i as many times as there are links on route i . We denote this vector by γ_{ISF} . In other words, we replace each throughput constraint by a set of constraints and the optimization problem (7) becomes

$$\begin{aligned} & \min_{\mathbf{d}} \mathbf{w}^T \mathbf{d} \\ & \mathbf{B}_1 \frac{\mathbf{r}}{\mathbf{d}(\mathbf{d} + \mathbf{r})} + \mathbf{B}_2 \frac{1}{(\mathbf{d} + \mathbf{r})} \leq \gamma_{\text{ISF}}^{-1}, \\ & \mathbf{d} > 0. \end{aligned} \quad (16)$$

Analogously to the store-and-forward case, the iteration step is now given by

$$\mathbf{u}_k = \mathbf{u}_{k-1} * \gamma_{\text{ISF}} * \left(\mathbf{B}_1 \frac{\mathbf{r}}{\mathbf{d}(\mathbf{d} + \mathbf{r})} + \mathbf{B}_2 \frac{1}{(\mathbf{d} + \mathbf{r})} \right), \quad (17)$$

where d_l is obtained from solving the infimum of the Lagrangian. In this case, d_l is the positive real-valued root x of the polynomial

$$w_l x^4 + 2r_l w_l x^3 + (r_l^2 w_l - \beta_l) x^2 - 2\alpha_l r_l x - \alpha_l r_l^2 = 0,$$

where α_l and β_l are the l th elements of $\mathbf{B}_1^T \mathbf{u}_{k-1}$ and $\mathbf{B}_2^T \mathbf{u}_{k-1}$, respectively. The scalability of this approach is not as good as it is with SF bound because of the increased number of constraints and additional computation related to the solution of \mathbf{d} .

VI. NUMERICAL EXAMPLES

In this section we consider two numerical examples to illustrate the methods and their properties.

A. Simple network example

We consider a network with 2 links and 2 traffic classes in a so called parking-lot configuration, see Figure 1. In this example, the load is $\rho = [5, 7]^T$ Mbps and the routing matrix is

$$\mathbf{A} = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}.$$




Fig. 1. Network with 2 links.

TABLE I
DIMENSIONING BASED ON AVERAGE THROUGHPUT WITH $\gamma_{\text{AVE}} = 1$ MBPS.

	c_1 (Mbps)	c_2 (Mbps)
BF	6.02	13.56
AVE	6.06	13.65

Thus, the link loads are $\mathbf{r} = [5, 12]^T$ Mbps. For this network explicit formulae are available for the per-class throughputs under balanced fairness [20]:

$$\begin{aligned} \Gamma_1^{\text{BF}}(\mathbf{c}) &= \left(\frac{1}{c_1 - r_1} + \frac{1}{c_2 - r_2} - \frac{1}{c_2 - r_1} \right)^{-1}, \\ \Gamma_2^{\text{BF}}(\mathbf{c}) &= c_2 - r_2. \end{aligned} \quad (18)$$

At the stability limit $\mathbf{c} = \mathbf{r}$, the flow-level performance tends to zero and some excess capacity is needed to achieve a given flow-level QoS criterion.

We first consider dimensioning based on average throughput such that the sum of the link capacities is minimized, i.e., the link costs are $w_1 = w_2 = 1$. The average throughput requirement is $\gamma_{\text{ave}} = 1$ Mbps. Since the exact throughputs are known for this network, we can compare the results from using (6) with the application of the exact throughput formulas (18). The exact solution can be obtained numerically by solving

$$\begin{aligned} & \min_{\mathbf{c}} \mathbf{w}^T \mathbf{c}, \\ & \left(\sum_i \frac{\rho_i}{\rho_{\text{tot}}} \frac{1}{\Gamma_i^{\text{BF}}(\mathbf{c})} \right)^{-1} \geq \gamma_{\text{ave}}, \\ & \mathbf{c} > \mathbf{r}. \end{aligned}$$

The resulting link capacities required to satisfy γ_{ave} are shown in Table I. Recall that the relation between the actual link capacity c_l and the excess capacity d_l resulting from (6) is simply, $c_l = r_l + d_l, \forall l$. In the table, the BF-row contains the results from applying the exact results under balanced fairness and the AVE row shows the results from (6), which is based on using the conservative SF bound for the throughput.

If the link capacities are fixed according to the BF-row in Table I the actual realized throughputs from (18) are 0.66 and 1.56 Mbps for class 1 and 2, respectively. The corresponding figures for the AVE-SF dimensioning are 0.70 and 1.65 Mbps.

To illustrate the per-class throughput methods, we let the target per-class throughput be the same for all classes so that $\gamma_i = 1$ Mbps, $i = 1, 2$. The results are shown in Table II, where LB refers to the lower bound on link capacities obtained from (8), BF to the exact result with (18), UB-SF to the upper bound on the link capacities from (13) and (14) and UB-ISF to the tighter upper bound from (17). The exact solution is

TABLE II
DIMENSIONING BASED ON PER-CLASS THROUGHPUT WITH $\gamma_i = 1$ MBPS,
 $i = 1, 2$.

	c_1 (Mbps)	c_2 (Mbps)
LB	6.00	13.00
BF	6.81	13.78
UB-ISF	6.87	13.86
UB-SF	7.00	14.00

TABLE III
PER-CLASS THROUGHPUTS GIVEN BY DIFFERENT DIMENSIONING
METHODS WITH $\gamma_i = 1$ MBPS, $i = 1, 2$.

	$\Gamma_1^{\text{BF}}(\mathbf{c})$ (Mbps)	$\Gamma_2^{\text{BF}}(\mathbf{c})$ (Mbps)
LB	0.53	1.00
BF	1.00	1.78
UB-ISF	1.04	1.86
UB-SF	1.13	2.00

obtained numerically by solving

$$\begin{aligned} \min_{\mathbf{c}} \mathbf{w}^T \mathbf{c}, \\ \Gamma_i^{\text{BF}}(\mathbf{c}) \geq \gamma_i, i = 1, 2, \\ \mathbf{c} > \mathbf{r}. \end{aligned}$$

The resulting per-class throughputs from (18) when using the capacities given by the different dimensioning methods are shown in Table III. Whereas the simpler methods, i.e., lower bound (8) and the upper bound (13), can be used to quickly bound the solution under balanced fairness, a more accurate estimate (and a tighter upper bound) of the required capacities is provided by the ISF method (17).

B. Large network example

The previous example illustrated the dimensioning methods with a simple network, where an explicit solution for the per-class throughput under balanced fairness is available. In general, explicit solutions for the throughput are not available. However, our dimensioning methods continue to be applicable in networks of arbitrary size.

We consider a larger network example with the flow classes (routes) as shown in Figure 2. The network consists of 2 gateway nodes and there are 18 flow classes (routes) in the network. We minimize the sum of the link capacities, i.e., the link costs $w_l = 1$, for all l , and we assume that the offered load $\rho_i = 1$ Mbps, for all i .

First we compare the average throughput dimensioning with $\gamma_{\text{ave}} = 0.1$ Mbit and per-class throughput dimensioning with $\gamma_i = 0.1$ Mbit for all i . To keep the presentation succinct we assume that the network graph is undirected thus resulting 12 links as numbered in Figure 2. The resulting link capacities are shown in Table IV. The labeling for the methods is as in Table II. As the throughput requirement is small relative to the load, all the methods yield comparable results.

The gap between the estimated upper and lower bounds on the link capacities depends on the throughput criterion. We illustrate this by evaluating the sum of the excess capacities,




Fig. 2. Example topology with two gateways and 18 flow classes (shown as arrows)

TABLE IV
LINK CAPACITIES FROM THE DIFFERENT DIMENSIONING METHODS FOR
THE LARGE NETWORK EXAMPLE.

Link	AVE	LB	UB-SF	UB-ISF
1	3.19	3.1	3.27	3.23
2	2.15	2.1	2.29	2.28
3	3.19	3.1	3.13	3.13
4	2.15	2.1	2.28	2.27
5	3.19	3.1	3.29	3.27
6	2.15	2.1	2.39	2.36
7	3.19	3.1	3.39	3.36
8	4.22	4.1	4.45	4.41
9	3.19	3.1	3.38	3.35
10	1.11	1.1	1.15	1.15
11	5.24	5.1	5.32	5.31
12	2.15	2.1	2.14	2.14

i.e., the value of the objective function, as a function of the per-class throughput requirement. In our example, the throughput requirement is again equal for all classes, $\gamma_i = \gamma$, for all i , and the traffic loads are $\rho_i = 1$ Mbps, for all i . The network links are now assumed to be unidirectional (22 links in total). The results are shown in Figure 3. When the throughput requirement γ is relatively high compared to the traffic loads, the accuracy of the SF method becomes worse suggesting the use of LB and ISF methods to characterize the link capacities. At small values of γ there is little difference between the SF and ISF methods. Recall that the actual required capacity under balanced fairness is between LB and ISF.

Finally, consider the case where the throughput target is kept fixed but the load is increased. In our example, $\gamma_i = 1$ Mbps, for all i , and the load is increased so that load is the same in all classes, i.e., $\rho_i = \rho$, for all i . As a function of the total offered load in the network, $\sum_i \rho_i$, we evaluate the sum of the normalized excess capacities $\sum_l \frac{d_l}{r_l}$, which characterizes the required amount of extra capacity relative to the link loads. The results are shown in Figure 4. For a fixed throughput requirement, when load is increased the required extra capacity relative to the link loads decreases sharply due to the multiplexing gain.




Fig. 3. Sum of excess link capacities, $\sum_l d_l$, as a function of γ . From top to bottom the curves are; SF method, ISF method and LB method.

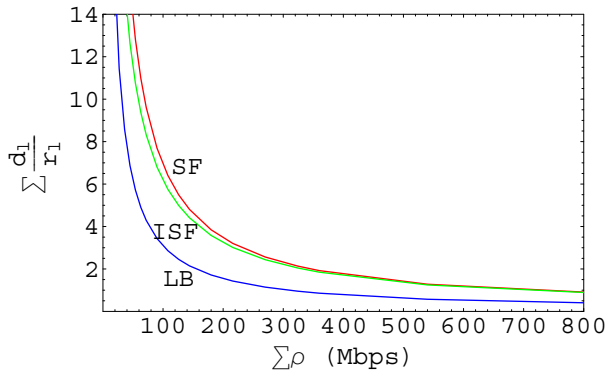


Fig. 4. Sum of normalized excess link capacities, $\sum_l d_l / r_l$, as a function of $\sum_i \rho_i$. From top to bottom the curves are; SF method, ISF method and LB method.

VII. CONCLUSIONS

We have provided methods for dimensioning of data networks with elastic traffic assuming that the bandwidth is shared according to balanced fairness. The average throughput dimensioning formulation consists of a single constraint and can be explicitly solved when applying the SF bound. The per-class throughput dimensioning has a constraint for each flow class. Using the upper bound for the per-class throughput we get a straightforward solution for the lower bound of the network cost. Using the lower bound formulas (SF and ISF) for the throughput we obtain upper bounds on the network cost. The associated optimization problems are more challenging. For this problem we developed an iteration scheme that is extremely simple to implement, especially in the case of the SF bound.

Our work is based on the balanced fairness assumption. While the bandwidth sharing in real networks does not exactly obey balanced fairness, it has been observed that balanced fairness also serves as a reasonable approximation to other sharing schemes, such as max-min or proportional fairness [5]. Therefore, we believe that the dimensioning methods provided here can be applied to easily obtain well-founded, although rough, estimates of the required link capacities to be used in

practical network planning.

The on-going work addresses two important extensions. The peak rates of the flows are often limited by, e.g., ADSL access rates, which has an impact also on dimensioning of the rest of the network. Another important issue is to include the QoS requirements of delay-sensitive traffic, such as VoIP traffic, into the model.

REFERENCES

- [1] A. K. Erlang, "Solution of some problems in the theory of probabilities of significance in automatic telephone exchanges," *Elektroteknikeren*, vol. 13, 1917.
- [2] D. Bertsekas and R. Gallager, *Data networks*, 2nd ed. Prentice-Hall, 1992.
- [3] F. Kelly, A. Maulloo, and D. Tan, "Rate control in communication networks: shadow prices, proportional fairness and stability," *Journal of the Operational Research Society*, vol. 49, pp. 237–252, 1998.
- [4] T. Bonald and A. Proutière, "Insensitive bandwidth sharing in data networks," *Queueing Systems*, vol. 44, pp. 69–100, 2003.
- [5] T. Bonald, L. Massoulié, A. Proutière, and J. Virtamo, "A queueing analysis of max-min fairness, proportional fairness and balanced fairness," *Queueing Systems*, vol. 53, pp. 65–84, 2006.
- [6] T. Bonald and A. Proutière, "On performance bounds for balanced fairness," *Performance Evaluation*, vol. 55, pp. 25–50, 2004.
- [7] T. Bonald, "Throughput performance in networks with linear constraints," in *Proceedings of 40th Annual Conference on Information Sciences and Systems (CISS)*, Mar. 2006.
- [8] M. Pioró, G. Malicskó, and G. Fodor, "Optimal link capacity dimensioning in proportionally fair networks," in *Proceedings of Networking 2002*, May 2002, pp. 277–288.
- [9] 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)*, Apr. 2002, pp. 2304 – 2311.
- [10] P. Nilsson and M. Pióro, "Solving dimensioning tasks for proportionally fair networks carrying elastic traffic," *Performance Evaluation*, vol. 49, pp. 371–386, 2002.
- [11] J. Charzinski, "Fun factor dimensioning for elastic traffic," in *Proceedings of ITC Specialist Seminar on Internet Traffic Measurement, Modeling and Management*, Sep. 2000.
- [12] J. Cohen, "The multitype phase service network with generalized processor sharing," *Acta Informatica*, vol. 12, pp. 245–284, 1979.
- [13] 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)*, Sep. 2000.
- [14] T. Bonald, P. Olivier, and J. Roberts, "Dimensioning high speed IP access networks," in *Proceedings of 18th International Teletraffic Congress*, Aug. 2003, pp. 241–251.
- [15] P. Lassila, A. Penttinen, and J. Virtamo, "Dimensioning of wireless mesh networks with flow-level QoS requirements," in *The 3rd ACM International Workshop on Performance Evaluation of Wireless Ad Hoc, Sensor, and Ubiquitous Networks (PE-WASUN 2006)*, Torremolinos, Spain, Oct. 2006.
- [16] T. Bonald and J. Virtamo, "Calculating the flow level performance of balanced fairness in tree networks," *Performance Evaluation*, vol. 58, pp. 1–14, 2004.
- [17] M. Jonckheere and J. Virtamo, "Optimal insensitive routing and bandwidth sharing in simple data networks," in *Proceedings of ACM SIGMETRICS*, Banff, Canada, Jun. 2005, pp. 193–204.
- [18] J. Leino and J. Virtamo, "Insensitive load balancing in data networks," *Computer Networks*, vol. 50, no. 8, pp. 1059–1068, 2006.
- [19] A. Penttinen, J. Virtamo, and R. Jäntti, "Performance analysis in multi-hop radio networks with balanced fair resource sharing," *Telecommunications Systems*, vol. 31, no. 4, pp. 315–336, Apr. 2006.
- [20] T. Bonald, A. Proutière, J. Roberts, and J. Virtamo, "Computational aspects of balanced fairness," in *Proceedings of 18th International Teletraffic Congress*, Aug. 2003, pp. 801–810.
- [21] M. Bazaraa, H. Sherali, and C. Shetty, *Nonlinear programming: theory and algorithms*, 2nd ed. Wiley, 1993.