

Optimal Partition of QoS Requirements with Discrete Cost Functions

Danny Raz Yuval Shavitt
{raz, shavitt}@research.bell-labs.com
Bell Laboratories, Lucent Technologies
101 Crawfords Corner Road
Holmdel, NJ 07733-3030

Abstract—The future Internet is expected to support applications with quality of service (QoS) requirements. For this end several mechanisms are suggested in the IETF to support signaling, the most promising among them is Diff-Serv. An important problem in this framework is how to partition the QoS requirements of an application along a selected path. The problem which is in general NP complete, was solved for continuous convex cost functions by Lorenz and Orda. This work concentrates on discrete cost functions, and presents efficient exact and approximated solutions for various conditions of the problem. We also show that the more complex problem of QoS sensitive routing with discrete cost functions is hard, but has a fully polynomial approximation scheme.

I. INTRODUCTION

The future networks are expected to support applications with quality of service (QoS) requirements. For this end, mechanisms are required to support signaling for connection establishment that include QoS routing and resource allocation. The QoS routing problem is to find a minimal cost path (or a multicast tree) in the network that can support the connection QoS requirements (such as delay). Along the selected path, resources (bandwidth, buffer space) should be optimally allocated to support the required QoS at a minimal cost. The latter can be formulated as an optimization problem for the partition of the end-to-end QoS requirements to local requirements along a path (or a multicast tree).

In general, the partition problem is intractable. The special case where the link cost functions, i.e., the function that describes the cost of allocating a QoS parameter on a link, are continuous convex cost functions was addressed recently by several works. Kodialam and Low [1] dealt with multicast trees for the strongly convex case. Lorenz and Orda [2] presented polynomial algorithms both for trees and paths for weakly convex cost functions and addressed the QoS routing problem [3].

This work concentrates on discrete cost functions, and presents efficient exact and approximated solutions for var-

ious cases. We first show that even the simplest possible discrete case, i.e., two level cost functions, is still intractable. We give an efficient dynamic programming solution for the special case where the QoS parameter domain is integer, but not necessarily convex. We present a sub-linear algorithm for the homogeneous convex case. Both solutions are demonstrated to be easily distributed with low communication and storage complexity. The same techniques are also used to establish similar algorithms for the multicast problem.

For the general discrete cost function case, we show a simple reduction of the QoS partition and the QoS routing problems to the restricted shortest path problem [4]. Using this reduction, one can easily derive an ϵ -approximation algorithm both for the QoS partition and routing problems in the unicast case. However, this reduction does not apply to the multicast case. Thus, we present a different fully polynomial approximation algorithm for the QoS partition problem that works both for the unicast and multicast case. Namely, we prove that for any approximation parameter ϵ our approximation algorithm finds a solution with cost not greater than $1 + \epsilon$ times the optimal cost, both for paths and trees. Moreover, we show that our approximation can solve also a more general class of non-discrete cost functions.

The discrete model used in this work lends itself more easily for practical purposes than its continuous counterpart. For example, in the Internet, *DiffServ* is suggested as a framework for QoS provisioning [5], [6]. In *DiffServ*, each packet can be classified to one of finitely many service classes. Such a class, for example, can guarantee a certain bound on the delay for passing through an AS (autonomous system), and is associated with a cost. Consider an application like IP telephony, that requires a delay bound of say 120mSec. To admit this call, we first have to select a route that can support the delay bound, but we would also like to pay as little as possible. Once a path is chosen, we still need to partition the delay bound requirement among the different ASs that comprise this path in a

way that results in a minimal cost. Each of the AS publishes the guaranteed delay and the cost for each service level, which, maps directly to our model.

The support of QoS has been the subject of excessive research. The specific aspect of resource allocation in this context has also been excessively studied, in particular, similar framework was studied by [7], [8], [2], [1]. The reader is referred to [9] for a survey on QoS multicast routing algorithms, though from a slightly different perspective.

The rest of the paper is organized as follows. In the next section we detail our model and define families of discrete cost functions. In Section III we prove that the QoS partition problem is NP-hard. The paper then focus on the unicast case: in Section IV, we solve the problem for the special case of integer functions; and in the next section we give an approximation algorithm for general discrete cost functions. In Section VI we extend these results to the multicast case. In Section VII we describe an approximation algorithm for the QoS routing problem, and in Section VIII we show that our approximation results also hold for non-discrete cost functions.

II. MODEL

A network is represented by a graph $G(V, E)$, each link $e \in E$ is associated with a discrete cost function $c_e : \mathcal{Q} \rightarrow \mathbf{R}$, that assigns a real positive value to each QoS parameter value. To simplify the discussion we sometime refer to the QoS parameter as delay.

In the unicast case, a path, \mathbf{p} , of length n between two end nodes is given, and that the QoS requirement is additive. Given a bound, \hat{Q} , on the end to end QoS requirement, the QoS partition problem is to find a vector $X = (x_1, \dots, x_n)$, s.t., $\sum_{i=1}^n x_i \leq \hat{Q}$, and $\sum_{i=1}^n c_i(x_i)$ is minimal. Note that, the case of bottleneck QoS requirement is trivial [2], and the multiplicative case can be easily reduced to the additive case by using the logarithm of the requirement [10], [2].

In the multicast case, a multicast tree, \mathbf{T} , with n nodes is given, and the QoS partition problem is to find a vector $X = (x_1, \dots, x_n)$, s.t., $\sum_{i \in \mathbf{p}} x_i \leq \hat{Q}$, for all paths, \mathbf{p} , in the tree and $\sum_{i=1}^n c_i(x_i)$ is minimal.

The QoS partition problem is called homogeneous if all the links have the same cost function.

A. Discrete cost functions

A general discrete cost function associates a cost with each discrete level of QoS. In the most general case, there may be infinitely many discrete QoS levels. We concentrate on the case where link i has i_k QoS levels, $q_{i_1}, \dots, q_{i_{i_k}}$. In such a case $c_i = (c_i(q_{i_1}), \dots, c_i(q_{i_{i_k}}))$.

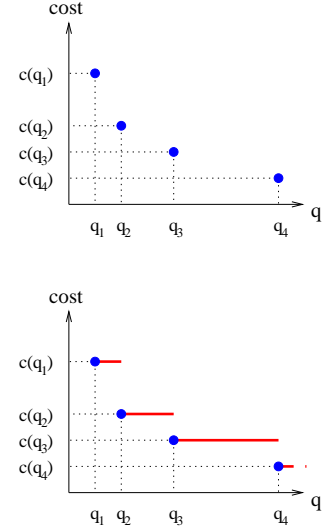


Fig. 1. A discrete cost function, and its representation as a step function.

Note that the representation of the discrete cost functions causes the input size to depend on the possible number of QoS levels, Q .

A convenient way to visualize a cost function is to consider a step function where the cost of a QoS parameter q is the cost of the biggest discretely defined QoS parameter $q_{i_j} \leq q$ (see Fig. 1). However, technically the function is defined only for the discrete points q_i , and it is easy to see that an optimum partition is always at these points, since sliding rightwards on a step increases the QoS parameter without decreasing the cost.

Next we define some special cases of cost functions

Definition 1: A cost function is called *integer* if it is defined only in a subset of the k points $1, \dots, k$.

Definition 2: A cost function is called *fully integer* if it is defined exactly at the k points $1, \dots, k$.

Note that by scaling, any discrete finite cost function (defined on the rationals) can be translated into an integer function. However, this may increase the cost of representing a set of functions exponentially, and thus translate a polynomial solution for an integer function to a pseudo-polynomial solution.

Definition 3: A cost function, c , is called *convex* if for every three points $q_i < q_j < q_l$ we have

$$c(q_j) \leq \frac{(q_j - q_i)c(q_i) + (q_l - q_j)c(q_l)}{q_l - q_i}$$

Definition 4: A cost function, c , is called *strongly convex* if for every three points $q_i < q_j < q_l$ we have

$$c(q_j) < \frac{(q_j - q_i)c(q_i) + (q_l - q_j)c(q_l)}{q_l - q_i}$$

The above definition requires that an intermediate point q_j is below the straight line connecting any two points one to its left (q_i) and one to its right (q_l).

III. HARDNESS RESULTS

In this section we prove that in general the QoS partition problem is NP-complete even if the discrete functions are convex. In particular, we show that even if the cost functions are the simplest non-trivial possible, containing only one-step functions, but different for every link, the problem is intractable.

Lemma 1: Let the cost function for link l be

$$c_l(i)^1 = \begin{cases} a_l & i < a_l \\ 0 & i \geq a_l \end{cases}$$

then determining whether the optimal solution to the QoS partition with limit \hat{Q} for a path of length n is \hat{Q} is equivalent to solving the *subset sum* problem [11, problem SP13] with a set of items a_1, \dots, a_n and a bound $B = \sum a_i - \hat{Q}$.

Proof: It is easy to show by comparing the problem definitions that the optimal cost of the above QoS partition with limit \hat{Q} for a path of length n is \hat{Q} if and only if there exist a subset $S \subset \{1, 2, \dots, n\}$, with $\sum_{i \in S} a_i = \hat{Q}$. \square

We will show in section V that although the problem is NP-complete, good approximation algorithms can be used to solve it.

IV. EXACT SOLUTIONS

In this section, we solve the QoS partition problem for integer cost functions. We first present a polynomial dynamic programming algorithm for the general case, and then give a sub-linear solution for the homogeneous case.

A. The general case

In this section we use dynamic programming to solve the QoS partition problem for a collection of integer cost functions. The only requirement is that all these functions can be defined on the same integer scale with no significant increase in their representation. We do not impose any other requirements on the functions, in particular, they need not be convex.

Let $cost(k, d)$ be the optimal cost of partitioning the QoS requirement d along the path l_1, \dots, l_k . Clearly, $cost(k, d)$ can be calculated by the following recursive formula

$$cost(k, d) = \min_{0 \leq i \leq d} cost(k-1, d-i) + c_k(i) \quad (1)$$

¹In our notation, c_l is defined in the points $q_1 = 0$, $q_2 = a_i$, and $c_l(0) = a_i$, $c_l(a_i) = 0$.

The minimal cost for the partition of requirement \hat{Q} along a path is thus given by calculating $cost(n, \hat{Q})$.

Theorem 1: The complexity of calculating the QoS partition in the case of general integer cost functions is $O(n\hat{Q}^2)$, and the memory requirement is $O(n\hat{Q})$.

Proof: For the calculation we need to keep a table of $cost(k, d)$ where $1 \leq k \leq n$ and $1 \leq d \leq \hat{Q}$. This requires a storage of $O(n\hat{Q})$ numbers. The calculation of each entry is done using Equation 1 which requires to minimize up to \hat{Q} sums, hence the calculation complexity. \square

B. Convex cost functions

For the case where the cost functions are fully integer and strongly convex one can apply the algorithm by Lorenz and Orda [2] to find a solution in $O(n \log n \hat{Q})$. Note that (strongly) convex functions have, at most, \hat{Q} different values for q_i , $1 \leq i \leq \hat{Q}$.

In this section, we consider the case where the cost functions are fully integer and convex, but require them all to be identical. We give an optimal algorithm with time complexity which is only $O(\log \hat{Q})$. To this end, we first prove the following lemmas

Lemma 2: The optimal QoS partition in the homogeneous fully integer strongly convex case results in all the QoS parameters taken from at most two successive values.

Proof: Suppose to the contrary that the lemma does not hold. Then the optimal partition contains, at least, two QoS values, $q_i < q_j$, s.t. $i + 1 < j$. Assume first that $j - i$ is even and let $q_m = (q_j - q_i)/2$ (the function $c(\cdot)$ is defined at q_m since it is fully integer). By Definition 4 we know that

$$\begin{aligned} c(q_m) &< \frac{(q_m - q_i)c(q_i) + (q_j - q_m)c(q_j)}{q_j - q_i} \\ &= \frac{c(q_i) + c(q_j)}{2} \end{aligned} \quad (2)$$

Thus a solution where the QoS parameter q_i and q_j are both replaced with q_m has a better cost while $2q_m = q_i + q_j$, which contradicts the optimality of the solution.

Assume now that $j - i$ is odd, and let $m = (j - i - 1)/2$. By applying definition 3 on the triplets q_i, q_m, q_{m+1} and q_m, q_{m+1}, q_j we get

$$\begin{aligned} c(q_m) + c(q_{m+1}) &< \\ (1 - \alpha)c(q_i) + \alpha c(q_{m+1}) + \alpha c(q_m) + (1 - \alpha)c(q_j) \end{aligned} \quad (3)$$

where $\alpha = 1/(q_m - q_i) = 1/(q_j - q_{m+1})$. Simple algebraic manipulations yield

$$c(q_m) + c(q_{m+1}) < c(q_i) + c(q_j) \quad (4)$$

Since $q_m + q_{m+1} = q_i + q_j$ the optimality assumption is contradicted for this case, too. \square

Corollary 1: In the optimal QoS partition in the homogeneous fully integer strongly convex case at least one link is allocated $\lfloor \hat{Q}/n \rfloor$.

In the same way we can prove the lemma for the weakly convex case:

Lemma 3: There exists an optimal QoS partition in the homogeneous fully integer convex case where all the QoS parameters are taken from at most two successive values.

Corollary 2: There exists an optimal QoS partition in the homogeneous fully integer convex case where at least one link is allocated $\lfloor \hat{Q}/n \rfloor$.

An optimal QoS partition is calculated as follows. $q_i = \lfloor \hat{Q}/n \rfloor$. The number of links that allocated q_i is given by finding the maximal x that solves the inequation $\hat{Q} \geq xq_i + (n-x)q_{i+1}$. Since $c(\cdot)$ is a fully integer function (assume normalized to the integers) we have $q_{i+1} = q_i + 1$ and thus $x = \lfloor n(q_i + 1) - \hat{Q} \rfloor$. $n - x$ links are allocated at q_{i+1} .

C. Distributed implementation

The dynamic program for the general integer cost function (section IV-A) can be easily distributed. Node i along the path can calculate $cost(i, d)$ for $1 \leq d \leq \hat{Q}$, based on the \hat{Q} values passed to it from node $i - 1$. When the calculation reaches the end node, it selects its optimal value, and passes back the optimal portion left for the path's prefix. This process continues until it reaches the originator. The total number of messages is only $2n$, while the bit complexity is $O(n\hat{Q})$. The storage requirement at each node is $O(\hat{Q})$. Note that the reservation is done in the reverse direction, thus the origin can start transmission after two way handshake. A formal description of the more complex multicast case is given in Section VI-B.

The homogeneous case of section IV-B can be calculated at the source node since the cost functions are known and equal. Once the two QoS parameter values that should be used are determined a reservation message with a counter stating how many reservations should be made for each value can propagate along the path towards the destination.

D. Discrete cost

All the results in this section can be applied to the dual case, where the cost is discrete and one wishes to find the best delay a certain cost can buy. This simple extension is omitted.




Fig. 2. The reduction of the QoS partition problem for general discrete cost functions to the restricted shortest path problem

V. APPROXIMATIONS

In this section we give a fully polynomial approximation scheme for the QoS partition problem with general discrete cost function. We assume here that the cost values are all integers. In this case we can rephrase the problem as follows.

Definition 5: [The QoS partition problem with general discrete cost function] Given n sets $S_i = \{a_{i1}, a_{i2}, \dots, a_{i|S_i|}\}$ of objects, with specific delays and costs, $delay(a_{ij}) = d_{ij} \in \mathbb{Z}^+$ and $cost(a_{ij}) = c_{ij} \in \mathbb{Z}^+$, and a delay bound $D \in \mathbb{Z}^+$, find a subset containing n objects, each from a different set, such that their total delay is bounded by D , and the total cost is minimized.

We assume that the cost function is non-increasing, i.e., for all i , $d_{ij} > d_{i,j'} \rightarrow c_{ij} \leq c_{i,j'}$. Denote by $c_{i_{max}}$, the maximum cost of any element in the set S_i , and by c_{max} the overall maximum cost. Clearly $\sum_i d_{i_{max}} \leq D$, otherwise there is no feasible solution. We denote by m the over all number of elements, that is: $m = \sum_{i=1}^n |S_i|$.

Our first observation is that there is a straight forward reduction from the QoS partition problem with general discrete cost function to the restricted shortest path problem [4]. Thus one can use Hassin's results to derive a fully polynomial approximation scheme.

Claim 1: Given an instance of the QoS partition (routing) problem with general discrete cost function one can construct a bi criteria² graph G' such that the cost of the restricted shortest path problem in G' equals to the cost of the QoS partition (routing) problem.

Proof: We replace the i th link by a set of l_i parallel links, each corresponds to a specific working point in the discrete cost function. More formally, given an instance of the QoS

²A graph where each edge is associated with both a cost and a delay.

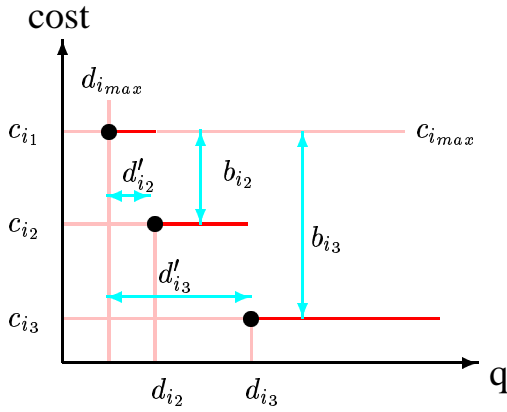


Fig. 3. The transition from the QoS partition problem with general discrete cost functions to the benefit discrete QoS partition problem

partition problem we build the graph G' with $n + 1$ nodes where nodes $i - 1$ and i are connected by l_i parallel links, with costs c_{i_j} and delays d_{i_j} (see Fig. 2). Since any simple path from node 0 to node n must choose exactly one of the edges between nodes $i - 1$ and i , a path with a delay bounded by D and cost C in G' defines a set with delay bounded by D and cost C in the QoS partition problem. \square

Note that the same reduction holds for the QoS routing problem with general discrete cost function, where the solution for the restricted shortest path problem determines both the links and the appropriate partition (see section VII).

There are two problems when applying Hassin's algorithm in this way. The first one is that the solution is complex and it is difficult to implement it. The more significant problem is that this result does not translate to multicast trees. Thus, the reduction does not hold for the QoS partition problems on trees. For this end we develop a different algorithm that can be generalized to multicast trees.

We begin by defining a variant of the problem, called the benefit discrete QoS partition problem, and proving that it has a fully polynomial approximation algorithm. We then show how to use this algorithm to achieve a fully polynomial approximation scheme for the QoS partition problem with general discrete cost function.

Definition 6: [The benefit discrete QoS partition problem] Given n sets $S_i = \{a_{i_1}, a_{i_2}, \dots, a_{i_{l_i}}\}$ of objects, with specific sizes and profits, $size(a_{i_j}) = d_{i_j} \in \mathbb{Z}^+$ and $profit(a_{i_j}) = b_{i_j} \in \mathbb{Z}^+$, and a delay bound $D \in \mathbb{Z}^+$, find a subset of at most n objects, each from a different set, such that their total size is bounded by D , and the total profit is maximized.

The main idea here is that the total profit for a given delay D represents the amount of cost one can save by

Pseudo Polynomial [PP] ($\{S_i\}_{i=1}^n, d_{i_j}, b_{i_j}, D$)

1. $b_{max} = \max_{i_j} b_{i_j}$.
2. for $p = 1$ to b_{max} :
3. $A(1, p) = \infty$; if $p = b_{i_1}$ then $A(1, p) = d_{i_1}$.
4. for $i = 2$ to n :
5. for $p = 1$ to b_{max}
6. $A(i, p) = \min\{\min_{1 \leq j \leq l_i} \{d_{i_j} + A(i - 1, p - b_{i_j})\}, A(i - 1, p)\}$.
7. return the largest p such that $A(n, p) \leq D$.

Fig. 4. Algorithm Pseudo Polynomial

Benefit Discrete [BD] ($\{S_i\}_{i=1}^n, d_{i_j}, b_{i_j}, \epsilon, D$)

1. $K = \frac{\epsilon b_{max}}{n}$
2. for each a_{i_j} : $b'_{i_j} = \lfloor \frac{b_{i_j}}{K} \rfloor$.
3. $B' = PP(\{S_i\}_{i=1}^n, d_{i_j}, b'_{i_j}, D)$.
4. if $(KB' < b_{max})$ output b_{max} , otherwise output KB' .

Fig. 5. Algorithm Benefit Discrete

allowing D more units of delay along the path, starting with any feasible (but maybe costly) initial solution (see Fig. 3). The objective is thus to gain as much savings as possible for every unit of delay.

First we show that the benefit discrete QoS partition problem has a pseudo polynomial algorithm, that uses dynamic programming. Then we use this algorithm in order to achieve the polynomial approximation scheme for the QoS partition problem.

Claim 2: Algorithm PP is a pseudo polynomial algorithm for the benefit discrete QoS partition problem, that works in $O((m + n)b_{max})$ time, where $b_{max} = \max_{i_j} b_{i_j}$.

Proof: Define $A(i, p)$, for $1 \leq i \leq n$, and $1 \leq p \leq n \times b_{max}$, to be the delay of the set with minimal delay that has at most i objects, each from a different set S_1, \dots, S_i , with benefit of exactly p . $A(i, p)$ is ∞ if no such set exists. Clearly

$$A(i, p) = \min\left\{\min_{1 \leq j \leq l_i} \{d_{i_j} + A(i - 1, p - b_{i_j})\}, A(i - 1, p)\right\}.$$

Now, the largest p such that $A(n, p)$ is smaller than D , is the optimal solution for the QoS partition problem. Since we need to compute $n \cdot b_{max}$ different values, and to compute $A(i, p)$ we need $l_i + 1$ steps, the over all complexity of the algorithm is $O((m + n)nb_{max})$. \square

Next we show how to use this algorithm in order to achieve an ϵ -approximation in polynomial time. We assume that the size of all elements is smaller than the bound D , since elements with bigger size cannot be used, and may as well be omitted.

Claim 3: Let B_{BD} be the profit outputted by Algorithm

Benefit Discrete (see Fig. 5). Then

$$B_{BD} \geq (1 - \epsilon)B_{opt}.$$

Proof: For every element a_{ij} , the profit considered by the algorithm may be smaller than the actual profit divided by K (as defined in line 1 of Fig. 5), but by no more than 1, i.e., $b'_{ij} \geq \frac{b_{ij}}{K} - 1$.

Let $profit_K(U)$ be the optimal solution for the instance with b' . Therefore for any set of elements U , $profit(U) - Kprofit_K(U) \leq nK$. The set S' computed by the PP algorithm must have at least the same profit as any other set, including the set computed by the optimal algorithm. Therefore

$$profit_{BD}(S') \geq Kprofit_K(U) \geq profit(U) - nK$$

$$profit_{BD}(S') \geq B_{opt} - \epsilon b_{max}.$$

Since the algorithm also considers (in the last step) the most profitable elements, $B_{BD} \geq b_{max}$. Therefore $profit_{BD}(S') \geq B_{opt} - \epsilon profit_{BD}(S')$, and,

$$B_{BD} \geq (1 - \epsilon)B_{opt}.$$

□

Since the running time of Algorithm Benefit Discrete is $O((n + m)n \lfloor \frac{b_{max}}{K} \rfloor) = O((n + m)n \lfloor \frac{n}{\epsilon} \rfloor)$, Theorem 2 follows.

Theorem 2: Algorithm Benefit Discrete is a fully polynomial approximation algorithm for the benefit discrete QoS partition problem.

Now we can describe the algorithm for the QoS partition problem with general discrete cost function. It works as follows (see Fig. 6). Given an approximation parameter ϵ , we construct a benefit discrete QoS partition problem that basically captures the amount of “saving” one can get starting from an obvious feasible solution. We find an ϵ/n -approximation to this value, and compute the resulting cost. We then iteratively remove the most expensive cost from one of the links and recompute the best cost. At the end we choose the best cost out of the (polynomially many) costs that we may have. The following claim follows immediately from the definitions.

Claim 4: For any feasible solution U to the benefit discrete QoS partition problem with b_{ij} , and d'_{ij} , there exists a set U' , which is a feasible solution for the original QoS partition problem with general discrete cost problem such that $profit(U) + cost(U') = \sum_i c_{imax}$.

This last claim proves that the algorithm finds a feasible set. We have to show both that the cost found by the algorithm is an ϵ -approximation of the optimal cost, and that the algorithm is polynomial.

General Discrete [GD] ($\{S_i\}_{i=1}^n, d_{ij}, c_{ij}, \epsilon, D$)

1. $C_{GD} = \sum_i c_{imax}$
2. repeat as long as none of the sets S_i is empty:
3. for each a_{ij} : $b_{ij} = c_{imax} - c_{ij}$; $d'_{ij} = d_{ij} - d_{imax}$.
4. $C'_{GC} = BD(\{S_i\}_{i=1}^n, d'_{ij}, b_{ij}, \epsilon/n, D - \sum_i d_{imax})$.
5. if $(\sum_i c_{imax} - C'_{GC} < C_{GD})$ then $C_{GD} = \sum_i c_{imax} - C'_{GC}$.
6. remove the element c_{max} from its set.

Fig. 6. Algorithm General Discrete

We already proved that the saving created by the algorithm, is almost as good as the saving created by any other algorithm. The problem is that this does not give a full proof for the approximation ratio, as the optimal cost may be much smaller than the total saving. However, if the optimal cost is at least $\sum_i c_{imax}/n$ we do have an ϵ -approximation since we used BD with $\epsilon' = \epsilon/n$. If the optimal cost is smaller, we will show that the biggest step in one of the links is not used in the optimal solution, hence we may delete it from the problem and start with a better upper bound. Thus, one of the iterations of the algorithm will find a solution which is an ϵ -approximation of the optimal cost. Since all sets are feasible, and we choose the one with minimal cost, the output of Algorithm General Discrete, is within an ϵ factor of the optimal cost.

For the formal proof we need the following two lemmas.

Lemma 4: If $\frac{C_{opt}}{\sum_i c_{imax}} \geq \frac{1}{k}$, and S' is a set with $profit(S') \geq (1 - \epsilon/k)b_{opt}$ for the benefit discrete QoS partition problem with $b_{ij} = c_{imax} - c_{ij}$ and $d'_{ij} = d_{ij} - d_{imax}$, then $cost(S') \leq (1 + \epsilon)c_{opt}$, for any constant $k > 1$.

Proof: By Claim 4, $cost(S') = \sum_i c_{imax} - profit(S')$, therefore,

$$cost(S') \leq \sum_i c_{imax} - (1 - \epsilon/k)b_{opt}.$$

Replacing again b_{opt} by $C_{opt} - \sum_i c_{imax}$ we get

$$cost(S') \leq \sum_i c_{imax} - (1 - \epsilon/k)(C_{opt} - \sum_i c_{imax}).$$

Using $\sum_i c_{imax} \leq kc_{opt}$ we get

$$cost(S') \leq \epsilon c_{opt} + c_{opt}(1 - \frac{\epsilon}{k}) \leq (1 + \epsilon)c_{opt}.$$

□

Lemma 5: If $\frac{C_{opt}}{\sum_i c_{imax}} < \frac{1}{n}$, then the element with the maximal cost cannot be in the set that achieves optimal cost.

Proof: Clearly,

$$C_{opt} < \frac{\sum_i c_{i_{max}}}{n} \leq c_{max}.$$

□

Now, Let O be the set that achieves an optimal cost. Let a_O be the element with the maximal cost in O . Since Algorithm GD deletes elements in order according to their cost values, there is an iteration in which a_O is the biggest element. By Lemma 5 in this iteration $\frac{C_{opt}}{\sum_i c_{i_{max}}} \geq \frac{1}{n}$ because the element with the maximal cost is in the optimal set. By Lemma 4, with $k = n$, the cost found in this iteration is bounded by $(1 + \epsilon)C_{opt}$. Since the algorithm chooses the best cost over all iteration its output is at least as good.

The running time of Algorithm General Discrete is bounded by $O(m^2 n \frac{n^2}{\epsilon})$ since we run the BD Algorithm at most m times, and the complexity of BD is $O(mn \lfloor \frac{b_{max}}{K} \rfloor) = O(mn \lfloor \frac{n}{\epsilon} \rfloor)$. However, we can replace the exhaustive search for the best $c_{i_{max}}$ by a binary search. This complicates the description of the algorithm but reduces the running time to $O(\frac{mn^3 \log m}{\epsilon})$. All together we have proven the following theorem.

Theorem 3: Algorithm GD is a fully polynomial approximation algorithm for the QoS partition problem with general discrete cost function.

VI. MULTICAST

A. Exact solutions

In this section we solve the QoS partitioning problem for a collection of integer cost functions in a multicast tree. As in the unicast case, the only requirement is that all these functions can be defined on the same integer scale with no significant increase in their representation. We do not impose any other requirements on the functions, in particular, they need not be convex. We begin by presenting a polynomial dynamic programming solution.

Let $l = (u, v)$ be a link in the multicast tree such that u is the parent of v , and let \mathcal{N}_l be the group of tree links connected to v except l , namely the group of links leading to v 's children. Remember that here n denotes the number of tree links. Let $cost(l, d)$ be the optimal cost of partitioning the QoS requirement d in the subtree of node v and the link l . Clearly, $cost(l, d)$ can be calculated by the following recursive formula

$$cost(l, d) = \min_{i \leq d} \sum_{e \in \mathcal{N}_l} cost(e, d - i) + c_l(i) \quad (5)$$

and the minimal cost for the partitioning of requirement \hat{Q} in the tree, \mathbf{T} is given by calculating $\sum_{e \in \mathcal{N}} cost(e, \hat{Q})$, where \mathcal{N} are the links emanating from the tree root.

Theorem 4: The cost of calculating the QoS partitioning in the case of general integer cost functions in a multicast tree is $O(n\hat{Q}^2)$, and memory requirement is $O(n\hat{Q})$.

Proof: For the calculation we need to keep a table of $cost(l, d)$ where $1 \leq l \leq n$ and $1 \leq d \leq \hat{Q}$. This requires a storage of $O(n\hat{Q})$ numbers. The calculation of each entry is done using Equation 5. For each entry the calculation cost is at most $\hat{Q}\mathcal{N}_l$, thus the overall calculation complexity is $\sum_{l \in \mathbf{T}} \hat{Q}\hat{Q}\mathcal{N}_l = n\hat{Q}^2$. □

For the homogeneous case where the cost functions are fully integer and convex, we can apply the calculation from section IV-B on the longest paths in the tree, remove them from the tree, and then repeat the process recursively on the remaining subtrees. This calculation is linear in the number of leaves.

B. Distributed implementation

The dynamic program for the general integer cost function (section VI-A) can be easily distributed. The root floods the tree with START messages. The START message carries in each link it traverses the cost function of the immediate upstream parent. A leaf that receives the START message from link l calculates its $cost(l, \cdot)$ entries and sends them on link l to its parent. A node that receives the $cost$ values from all its children, calculates its \hat{Q} entries and sends them to its parent. When the root receives the $cost$ calculations from all its children it initiates the reservation phase by sending a BUDGET message. This message carries the QoS remaining budget downstream. A node that receives the BUDGET message uses the appropriate entry in the cost table it calculated before locating the QoS parameter allocation in its upstream link. It asks its upstream neighbor to allocate this amount using the RESERVE message, and sends a BUDGET message with the remainder of the budget downstream.

The total number of messages is only $4n$, while the bit complexity is $O(n\hat{Q})$. The storage requirement for each link is $O(\hat{Q})$, thus, for a node with \mathcal{N} children the storage requirement is $\mathcal{N}\hat{Q}$. The time complexity is $O(3h)$ where h is the tree height.

Figures 7 and 8 give a formal description of the algorithm. Each node has the following variables: p the index of the parent link; $cost(l, \cdot)$ a vector with the optimal cost calculated by the node downstream link l ; $cost(\cdot)$ a vector with the results of the local optimal cost calculation; $value(\cdot)$ a vector with the QoS parameter that gives the best cost; and $gotit(\cdot)$ a binary vector to monitor the receipt of $cost$ calculations from the child links. In addition, a node holds a discrete cost function, $c_l(\cdot)$, for each of its downstream links l , and a list of its downstream links \mathcal{N} .

Distributed Multicast Partitioning

1. For $\text{START}(c_l(\cdot))$ from link l
2. $p \leftarrow l$
3. $c_p(\cdot) \leftarrow c_l(\cdot)$
4. foreach $i \in \mathcal{N}_l$
5. sent $\text{START}(c_i(\cdot))$ on link i
6. $\text{gotit}(i) \leftarrow \text{false}$
7. if $|\mathcal{N}_p| = 0$
8. foreach $j < \hat{Q}$
9. $\text{cost}(j) = \min_{i < j} c_l(j)$
10. $\text{value}(j) = \arg \min_{i < j} c_l(j)$
11. send $\text{cost}(l, \cdot)$ on link l
12. For $\text{cost}(\cdot)$ from link l
13. $\text{cost}(l, \cdot) \leftarrow \text{cost}(\cdot)$
14. $\text{gotit}(l) \leftarrow \text{true}$
15. if $\forall i \in \mathcal{N}_p : \text{gotit}(i) = \text{true}$
16. foreach $j < \hat{Q}$
17. $\text{cost}(j) = \min_{i < j} \sum_{e \in \mathcal{N}_p} \text{cost}(e, d - i) + c_p(i)$
18. $\text{value}(j) = \arg \min_{i < j} \sum_{e \in \mathcal{N}_p} \text{cost}(e, d - i) + c_p(i)$
19. send $\text{cost}(\cdot)$ on link p
20. For $\text{BUDGET}(d)$
21. send $\text{RESERVE}(\text{value}(d))$ on link p
22. foreach $i \in \mathcal{N}_l$
23. send $\text{BUDGET}(d - c_p(\text{value}(d)))$ on link i
24. For $\text{RESERVE}(d)$ from link l
25. reserve d delay guarantee on link l

Fig. 7. Distributed Multicast Partitioning — a non-root node algorithm

Distributed Multicast Partitioning

1. For a request with demand \hat{Q}
2. foreach $i \in \mathcal{N}$
3. sent $\text{START}(c_i(\cdot))$ on link i
4. $\text{gotit}(i) \leftarrow \text{false}$
5. For $\text{cost}(\cdot)$ from link l
6. $\text{cost}(l, \hat{Q}) \leftarrow \text{cost}(\hat{Q})$
7. $\text{gotit}(l) \leftarrow \text{true}$
8. if $\forall i \in \mathcal{N} : \text{gotit}(i) = \text{true}$
9. $\text{cost} = \sum_{e \in \mathcal{N}} \text{cost}(e, \hat{Q})$
10. foreach $i \in \mathcal{N}$
11. send $\text{BUDGET}(\hat{Q})$ on link i
12. For $\text{RESERVE}(d)$ from link l
13. reserve d delay guarantee on link l

Fig. 8. Distributed Multicast Partitioning — the root algorithm

C. Approximations

As mentioned before, in the multicast trees case we cannot use the reduction of Claim 1, since it will require finding a restricted shortest tree rather than a restricted shortest path.³

Thus, in order to derive a fully polynomial approximation scheme, we follow the steps we took in the path case: starting from a feasible (but possibly costly) partition on the tree, we try to use the extra delay we have in order to save as much cost as possible. We use many of the notations from Section V, and we assume that the multicast tree is a binary tree. It is easy to verify that for any multicast tree, \mathbf{T} , with n nodes, there exists a binary multicast tree \mathbf{T}' , with at most $2n$ nodes, with exactly the same optimal cost.⁴ The amount of saving is expressed in the following definition.

Definition 7: [The benefit discrete QoS partition problem on trees] Given a tree \mathbf{T} , with n nodes, sets $S_i = \{a_{i_1}, a_{i_2}, \dots, a_{i_{l_i}}\}$ of objects, with specific sizes and profits, $\text{size}(a_{i_j}) = d_{i_j} \in \mathbb{Z}^+$ and $\text{profit}(a_{i_j}) = b_{i_j} \in \mathbb{Z}^+$, and a delay bound $D \in \mathbb{Z}^+$, find a subset of at most n objects, $a_{i_{j_i}}$, each from a different set, such that $\sum_{i \in \mathbf{p}} d_{i_{j_i}} \leq D$, for all paths, $\mathbf{p} \in \mathbf{T}$, and $\sum_{i=1} b_{i_{j_i}}$ is maximized.

Next we prove that the benefit discrete QoS partition problem on trees has a pseudo polynomial algorithm, that uses dynamic programming. An algorithm similar to the dynamic programming algorithm from Section VI-A is not good enough because it is polynomial in the delay bound D and not the maximal cost c_{max} . In the multicast tree case the rules of the delay and the cost (benefit) is not symmetric since the cost (benefit) is computed over the entire tree while the delay bound is true for any path in the tree. Recall that we assume that the multicast tree \mathbf{T} is a binary tree. We also assume that node n is the root of the tree, and that i_L and i_R are the indices of the left and right children of node i .

Claim 5: Algorithm PPM is a pseudo polynomial algorithm for the benefit discrete QoS partition problem on multicast trees, that works in $O((m+n)nb_{max}^2)$ time, where $b_{max} = \max_{i,j} b_{i_j}$.

Proof: Define $A(i, p)$, for $1 \leq i \leq n$, and $1 \leq p \leq n \cdot b_{max}$, to be d iff there exists a QoS partition of the subtree rooted at node i with minimal delay bound d and profit p , and no other partition of this subtree with benefit p

³To the best of our knowledge approximating a restricted shortest tree for graphs with costs and delays, is an open problem.

⁴This is done by replacing the outgoing edges at any node that has more than two children, by a binary tree in which these children are the leaves, and all the links to internal nodes have both zero cost and zero delay.

| Pseudo | Polynomial | Multicast | [PPM] |
|--------|------------|-----------|---|
| | | | $(\mathbf{T}, \{S_i\}_{i=1}^n, d_{ij}, b_{ij}, D)$ |
| 1. | | | $b_{max} = \max_{ij} b_{ij}.$ |
| 2. | | | for $p = 1$ to b_{max} : |
| 3. | | | for all leaves l : |
| 4. | | | $A(l, p) = 0.$ |
| 5. | | | for $i = 2$ to n : |
| 6. | | | for $p = 1$ to b_{max} |
| 7. | | | $A(i, p) = \min_{0 \leq l \leq p} \{ \max \{ \min \{ A(i_L, l), \min_{a \in S_{i_L}} A(i_L, l - profit(a)) + d(a) \}, \min \{ A(i_R, p - l), \min_{a' \in S_{i_R}} A(i_R, p - l - profit(a')) + d(a') \} \} \}$ |
| 8. | | | return the largest p such that $A(n, p) \leq D.$ |

Fig. 9. Algorithm Pseudo Polynomial Multicast

| General Discrete for Multicast trees | [GDM] |
|--------------------------------------|--|
| | $(\{S_i\}_{i=1}^n, d_{ij}, c_{ij}, \epsilon, D)$ |
| 1. | $C_{GD} = \sum_i c_{i_{max}}$ |
| 2. | repeat as long as none of the sets S_i is empty: |
| 3. | for each a_{ij} : $b_{ij} = c_{i_{max}} - c_{ij}$; $d'_{ij} = d_{ij} - d_{i_{max}}$. |
| 4. | $C'_{GC} = BDM(\{S_i\}_{i=1}^n, d'_{ij}, b_{ij}, \epsilon/n, D - \sum_i d_{i_{max}}).$ |
| 5. | if $(\sum_i c_{i_{max}} - C'_{GC} < C_{GD})$ then $C_{GD} = \sum_i c_{i_{max}} - C'_{GC}.$ |
| 6. | remove the element $c_{i_{max}}$ from its set. |

Fig. 10. Algorithm General Discrete for multicast trees

has smaller delay bound. $A(i, p)$ is ∞ if no such partition exists. At node i , one can choose the amount of profit and delay from each of the sets associated with the left and the right children of i . One can also choose not to choose each of these elements. In any case, the relation defined by line 7 of Fig. 9 holds. Therefore the value of $A(n, p)$, is the best possible delay bound for the given tree \mathbf{T} with profit p , and the algorithm outputs the optimal value.

Since we need to compute $n \cdot nb_{max}$ different values, and to compute $A(i, p)$ we need at most $b_{max}(l_i + 1)$ steps, the over all complexity of the algorithm is $O((m+n)nb_{max}^2)$. \square

We apply now the rounding technique from Section V to achieve an approximation scheme for this case. We use Algorithm BD with a call to PPM rather than PP (this version is called BDM⁵). This results in an ϵ -approximation algorithm that runs in time $O((m+n)n(\lfloor \frac{b_{max}}{K} \rfloor)^2) = O(\frac{mn^3}{\epsilon^2})$.

Fig. 10 presents a slightly modified version of the GD Algorithms for the multicast trees case. Applying Claim 4 and Lemmas 4, 5 for GMD we get the following theorem.

Theorem 5: Algorithm General Discrete Multicast is a fully polynomial approximation algorithm for QoS parti-

⁵The formal description of BDM is omitted.

tion problem on multicast trees with general discrete cost function. It has time complexity of $O(\frac{mn^4 \log m}{\epsilon^2})$.

VII. THE QoS REQUIREMENTS ROUTING PROBLEM

In this section we formally define the QoS requirements routing problem with discrete cost functions. We then show that it is an NP hard problem and show how to obtain a fully polynomial approximation scheme for this problem.

Definition 8: [The QoS requirements routing problem with discrete cost functions] Given a graph $G = (V, E)$, each edge is associated with a discrete cost function, and a delay bound D , find a path \mathbf{p} , and a partition such that the delay along the chosen path is bounded by D , and the cost along this path is the minimal possible.

Clearly this problem is NP-complete since the QoS partition problem with general discrete cost function is a special case of it (when the graph G is a line). Even the simpler problem of finding the best path without specifying the QoS partition along it is NP-complete, this follows from a straight forward reduction to the restricted shortest path problem [4].

Claim 1 provides a constructive way to construct an instance of the restricted shortest path problem from a given instance of the QoS routing problem with general discrete cost function. Using it, and the approximation results from [4], we get the following Theorem.

Theorem 6: The QoS requirements routing problem with discrete cost functions has a fully polynomial approximation scheme.

VIII. GENERAL COST FUNCTIONS

In this section we discuss general cost functions. That is, we do not assume any properties such as monotonicity or convexity. Such a function, which is defined for all cost values, is represented by a constant number of bytes. It computes the delay guarantee for a given cost in polynomial time in the representation of the input number. Note that although the discrete functions defined in Section II look similar to the general cost functions, the representation of such a function is linear in the number of points and not constant. For example, the function $cost(d) = \frac{1}{d}$ is defined for every $d > 0$, but if we want to represent it by a discrete function we would have to explicitly give the cost value of each possible delay. To emphasize this point we use here $f_i : Cost \rightarrow Delay$ for the reverse function $cost^{-1}$, thus for the above example $f_i^{-1}(d) = \frac{1}{d}$. For such functions we can prove the following theorem.

Theorem 7: For a set of n functions $f_i : Cost \rightarrow Delay$, such that for each such a function either f_i^{-1} , or a pair $(c_i, d_i = f(c_i))$, with $d_i < D/n$ is given, a graph G ,

a delay bound D , and an approximation parameter ϵ , one can:

1. Given a path i, j , find an ϵ -approximation algorithm for the QoS partition problem.
2. Given a multicast tree, \mathbf{T} , find an ϵ -approximation for the QoS partition problem on multicast trees.

Proof: First observe that given the reverse function f_i^{-1} one can compute $f_i^{-1}(D/n) = c_i$, and use the pair $(c_i, D/n)$. These pairs define a feasible solution, both for the path and multicast tree problems with cost bounded by $c_{max} = \sum_{i=1}^n c_i$. Therefore no optimal solution will use any cost bigger than c_{max} . Observe now that the pseudo polynomial algorithms PP and PPM work also if each step function has c_{max} steps. For each $A(i, p)$ we have to check up to $c_{max} + 1$ different values, and therefore the overall complexity of the pseudo polynomial algorithms will increase by a factor of c_{max}/m . We can use the algorithms from Section V and Section VI-C to derive fully polynomial approximation schemes for the QoS partition problem and the QoS partition problem on multicast trees, respectively.

□

A similar proof can be used to prove an ϵ -approximation scheme for the routing problem with general cost functions. However it requires either to generalize our results from Section V to the QoS routing problem, or to show that Hassin's algorithm can be used. This is beyond the scope of this paper.

IX. DISCUSSION

In this paper we studied QoS partition and routing problems. We concentrated on discrete cost functions, that are both theoreticly interesting and have practical applications in IP networks. An interesting direction, with possible practical significant, is studying a distributed implementation of the approximation algorithms presented here (in the same spirits of the results of Sections IV-C and VI-B). The main open problem, however, remains the multicast trees routing problem.

REFERENCES

- [1] Murali Kodialam and Steven H. Low. Resource allocation in a multicast tree. In *IEEE INFOCOM'99*, pages 262 – 266, March 1999.
- [2] Dean H. Lorenz and Ariel Orda. Optimal partition of QoS requirements on unicast paths and multicast trees. In *IEEE INFOCOM'99*, pages 246 – 253, March 1999.
- [3] Dean H. Lorenz and Ariel Orda. QoS routing on networks with uncertain parameters. *IEEE/ACM Transactions on Networking*, 6(6):768 – 778, December 1998.
- [4] Refael Hassin. Approximation schemes for the restricted shortest path problem. *Mathematics of Operations Research*, 17(1):36 – 42, February 1992.

- [5] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss. An architecture for differentiated services. Internet RFC 2475, November 1998.
- [6] Walter Weiss. QoS with differentiated services. *Bell Labs Technical Journal*, 3(4):48–62, oct – dec 1998.
- [7] R. Nagarajan, J.F. Kurose, and D. Towsley. Heirat: The heidelberg resource administration technique design philosophy and goals. In *IFIP Workshop on the Performance Analysis of ATM Systems*, January 1993.
- [8] V. Firoiu and D. Towsley. Call admission and resource reservation for multicast sessions. In *IEEE INFOCOM'96*, March 1996.
- [9] Shigang Chen and Klara Nahrstedt. An overview of quality-of-service routing for the next generation high-speed networks: Problems and solutions. *IEEE Network Magazine*, 12(6), Nov./Dec. 1998.
- [10] D. Bertsekas and R. Gallager. *Data Networks*. Prentice Hall, second edition, 1992.
- [11] Michael R. Garey and David S. Johnson. *Computer & Intractability: A Guide to the Theory of NP-Completeness*. W H Freeman, November 1979.