

near the edges, where link bandwidth is limited. Perhaps wireless links, with their limited capacity, will be a major reason to implement multiple packet classes. Some wireless links, and some data streams are time varying with large time scales. Such time variations, coupled with the limitations imposed by large round trip times in congestion control loops, may forever insure that queuing delay, along with bandwidth, has to be explicitly addressed, pointing to a need for multiple packet classes.

Our goal for fairness is the following: provide at least the same level of satisfaction to throughput sensitive users when there are two packet classes as when there was only one class. An interesting question is how we can accommodate a multidimensional QoS profile within the notions of network-wide fairness. Another question is how to implement congestion pricing for multiple class networks. Progress in this direction was recently reported in [5] and [11].

In practice, the expectations and requirements of network users almost always involve delay and loss, whether or not explicitly stated in a service level specification or agreement. Perhaps increasing network resources in a timely fashion, relying on statistical multiplexing, and controlling admission to networks will someday make throughput the only relevant QoS measure. However, if some users are more tolerant to delay and loss than others, and if burstiness of aggregate traffic streams in some links cannot be avoided, the use of multiple packet classes can typically improve the margin of protection against unexpected or unavoidable stresses on the network. Similar conclusions are reached in [2].

REFERENCES

- [1] S. Athuraliya, D. Lapsley, and S. Low, "An enhanced random early marking algorithm for Internet flow control," in *Proc. INFOCOM'2000*, Israel, 2000.
- [2] S. Bajaj, L. Breslau, and S. Shenker, "Is service priority useful in networks?," in *Proc. ACM Sigmetrics'98*, Madison, WI, June 1998, pp. 66–77.
- [3] D. D. Clark, "Adding service discrimination to the Internet," *Telecommun. Policy*, vol. 20, pp. 33–37, 1996.
- [4] R. J. Gibbens and F. P. Kelly, "Resource pricing and the evolution of congestion control," *Automatica*, vol. 35, 1999. 1969–1985.
- [5] —, "A note on packet marking at priority queues," *IEEE Trans. Automat. Contr.*, vol. 47, pp. 1016–1020, June 2002.
- [6] F. P. Kelly, A. K. Maulloo, and D. K. H. Tan, "Rate control in communication networks: Shadow prices, proportional fairness and stability," *J. Oper. Res. Soc.*, vol. 49, pp. 237–252, 1997.
- [7] F. P. Kelly, "Models for a self-managed Internet," *Phil. Trans. Royal Soc. A358*, pp. 2335–2348, 2000.
- [8] —, "Mathematical modeling of the Internet," in *Mathematics Unlimited—2001 and Beyond*. Berlin, Germany, 2001, pp. 685–702.
- [9] P. B. Key and D. R. McAuley, "Differential QoS and pricing in networks: Where flow control meets game theory," in *Proc. Inst. Elect. Eng. Software*, vol. 146, 1999, pp. 39–43.
- [10] P. Key, D. McAuley, P. Barham, and K. Laevens, "Congestion pricing for congestion avoidance," Microsoft Research, Cambridge, U.K., Tech. Rep. MSR-TR-99-15, 1999.
- [11] P. Key, L. Massoulié, and J. Shapiro, "Service differentiation for delay-sensitive applications: An optimization-based approach," Cambridge, U.K., Tech. Rep. MSR-TR-2001-115, Nov. 2001.
- [12] L. Massoulié and J. Roberts, "Bandwidth sharing: Objectives and algorithms," *Proc. IEEE INFOCOM'99*, pp. 1395–1403, Mar. 1999.
- [13] D. Wischik, "How to mark fairly," presented at the M.I.T. Workshop Internet Service Quality Economics, Cambridge, MA, 1999.

On Localized Control in QoS Routing

Srihari Nelakuditi, Srivatsan Varadarajan, and Zhi-Li Zhang

Abstract—In this note, we study several issues in the design of *localized quality-of-service (QoS) routing schemes that make routing decisions based on locally collected QoS state information (i.e., there is no network-wide information exchange among routers)*. In particular, we investigate the *granularity of local QoS state information and its impact on the design of localized QoS routing schemes from a theoretical perspective*. We develop two theoretical models for studying localized proportional routing: one using the link-level information and the other using path-level information. We compare the performance of these localized proportional routing models with that of a global optimal proportional model that has knowledge of the global network QoS state. We demonstrate that using only coarser-grain path-level information it is possible to obtain near-optimal proportions. We then discuss the issues involved in implementation of localized proportional routing and present some practical schemes that are simple and easy to implement.

Index Terms—Localized proportional routing, quality-of-service (QoS) routing.

I. INTRODUCTION

In quality-of-service (QoS)-based routing [2], [6], [23], paths for flows are selected based upon knowledge of the resource availability (referred to as *QoS state*) at network nodes (i.e., routers) and the QoS requirements of the flows. This knowledge, for example, can be obtained through (periodic) information exchange among routers in a network. Under this approach, which we refer to as the *global QoS routing approach*, each router constructs a global view of the network QoS state by piecing together the QoS state information obtained from other routers, and performs path selection based on this global view of the network state. Examples of the global QoS routing approach are various QoS routing schemes [4], [23] based on QoS extensions to the OSPF routing protocol as well as the ATM PNNI routing protocol. Global QoS routing schemes work well when each source node has a reasonably *accurate* view of the network QoS state. However, as the network resource availability changes with each flow arrival and departure, maintaining an accurate network QoS state is impractical, due to the prohibitive communication and processing overheads entailed by frequent QoS state information exchange. In the presence of *inaccurate* information regarding the network QoS state, global QoS routing schemes may suffer degraded performance as well as potential instability [22], [14].

As a viable alternative to the global QoS routing approach, in [15] we proposed a novel *localized* approach to QoS routing. Under this localized QoS routing approach, instead of (periodically) exchanging information with other routers to obtain a global view of the network QoS state, a *source* router attempts to *infer* the network QoS state from *locally collected flow statistics* such as flow arrival/departure rates and flow blocking probabilities, and performs path selection based on this local information. As a result, the localized QoS routing approach avoids the drawbacks of the conventional global QoS routing approach such as degraded performance in the presence of inaccurate routing information. Furthermore, it has several important advantages: *minimal* communication overhead, *no* processing overhead at *core* routers, and *easy* deployability.

Manuscript received October 19, 2000; revised June 30, 2001 and November 25, 2001. Recommended by the Guest Editors M.-B. Gong and T. Başar.

The authors are with the Department of Computer Science and Engineering, University of Minnesota, Minneapolis, MN 55455 USA.

Publisher Item Identifier S 0018-9286(02)05508-3.

In this note, we investigate an important and fundamental issue in the design of localized QoS routing schemes—the *granularity* of locally collected QoS state information and its impact on the convergence process of these schemes and their performance. We consider flow statistics collected at two different granularity levels: *link* level and *path* level. At the (*finer*) link level, a source node collects *both* the blocking statistics (i.e., whether a flow is blocked or not) of flows routed along a path from the source node to a destination node *and*, in the case of a blocked flow, the identity of the link where the flow is blocked. The latter information can be gathered, for example, by attaching the identity of the link in the flow setup failure notification sent back to the source node. At the (*coarser*) path level, a source node collects *only* the flow blocking statistics for each path between the source node and a destination node. Clearly, the path-level flow statistics are easier to collect and maintain, but they also convey less precise information regarding the (global) network QoS state.

Using the flow blocking statistics collected at link and path levels, in this note we propose theoretical models to study the impact of granularity. These models are developed based on the notion of *virtual capacity* of a link or a path as *perceived* by a source node. The virtual capacity of a link or a path is computed as a function of the amount of offered load and the corresponding observed blocking probability on that link or path. Through numerical investigation, we show that it is possible to design localized proportional routing schemes that converge to a stable point. We find that though granularity of information does have impact on the rate of convergence and the equilibrium blocking probability, the performance penalty due to coarser path-level information is not significant. Based on these theoretical results, we proceed to develop practical localized proportioning strategies that are simple and easy to implement.

The remainder of the note is organized as follows. In Section II, we introduce the notion of virtual capacity as well as the (ideal) global optimal proportional QoS routing model. In Section III, we present the two theoretical localized QoS routing models, one using the link-level flow blocking statistics, and the other using the path-level flow blocking statistics. The issues in practical implementation of localized routing schemes are discussed in Section IV. We conclude the note in Section V.

II. PROPORTIONAL ROUTING: GLOBAL VERSUS LOCAL

In this section, we first lay out the basic assumptions regarding the *proportional* QoS routing models we consider in this note. We then present an ideal global proportional QoS routing model, where we assume that the *complete topology* information of the network as well as the *offered traffic load* between every source destination pair are *known*. The performance of the global optimal proportional routing scheme will serve as the basis for comparing the performance of localized proportional QoS routing schemes. Then, we introduce the notion of virtual capacity and briefly describe how it can be used in designing localized proportional routing schemes.

In all the QoS routing models we consider in this note, we assume that *source routing* (also referred to as *explicit routing*) is used. More specifically, we assume that the network topology information is available to all source nodes (e.g., via the OSPF protocol), and one or multiple *explicit-routed* paths are set up *a priori* for each source and destination pair using, e.g., MPLS [1]. Flows arriving at a source to a destination are routed along one of the explicit-routed paths (hereafter referred to as the *candidate* paths). For simplicity, we assume that all flows have the same bandwidth requirement—one unit of bandwidth.¹ When a flow is routed to a path where one or more of the constituent

links have no bandwidth left, this flow will be blocked. The performance metric in our study will be the overall blocking probability experienced by flows. We assume that flows from a source to a destination arrive randomly with a Poisson distribution, and their holding time is exponentially distributed. Hence, the *offered* traffic load between a source–destination pair can be measured as the product of the average flow arrival rate and holding time. Given the offered traffic load from a source to a destination, the task of proportional routing is to determine how to distribute the load (i.e., route the flows) among the candidate paths so as to minimize the overall blocking probability experienced by the flows.

A. Global Optimal Proportional Routing

Global optimal proportional routing problem has been studied extensively in the literature (see [21] and references therein). Given the global knowledge of the network topology and offered traffic loads, the *optimal proportions*, for distributing flows among the candidate paths between each source–destination pair, can be computed as follows. Consider an arbitrary network topology with N nodes and L links. For $l = 1, 2, \dots, L$, the capacity of link l is $c_l > 0$, which is assumed to be fixed and known. Let $\sigma = (s, d)$ denote a source–destination pair in the network. Let λ_σ denote the average arrival rate of flows arriving at source node s destined for node d . The average holding time of the flows is μ_σ . Recall that each flow is assumed to request one unit of bandwidth, and that the flow arrivals are Poisson, and flow holding times are exponentially distributed. Thus the offered load between the source–destination pair σ is $\nu_\sigma = \lambda_\sigma / \mu_\sigma$. Let R_σ denote the set of candidate paths for routing flows between the source–destination pair σ . The global optimal proportional problem can be formulated [8]–[10] as the problem of finding the optimal proportions $\{\alpha_r^*, r \in R_\sigma\}$ (where $\sum_{r \in R_\sigma} \alpha_r^* = 1$) for each source–destination pair σ , such that the overall flow blocking probability in the network is minimized. Or equivalently, finding the optimal proportions $\{\alpha_r^*, r \in R_\sigma\}$ such that the total carried traffic in the network

$$W = \sum_{\sigma} \sum_{r \in R_\sigma} \alpha_r \nu_\sigma (1 - b_r) \quad (1)$$

is maximized.

In (1), b_r denotes the blocking probability along path r . Under the *link independence* assumption [8]–[10], b_r can be expressed as follows:

$$b_r = 1 - \prod_{l \in r} (1 - b_l) \quad (2)$$

where $l \in r$ means that link l is part of route r , and b_l is the blocking probability of link l . The blocking probability b_l of link l is in turn given by the Erlang loss formula

$$b_l = E(\nu_l, c_l) = \frac{\nu_l^{c_l}}{c_l!} \sum_{n=0}^{c_l} \frac{\nu_l^n}{n!} \quad (3)$$

Here, the load offered on link l , ν_l is the sum of all the *reduced* loads (after independent load *thinning*) from any source–destination pair σ which has a route passing through link l , i.e.,

$$\nu_l = \sum_{\sigma} \sum_{r \in R_\sigma: l \in r} \alpha_r^* \nu_\sigma \prod_{m \in r - \{l\}} (1 - b_m). \quad (4)$$

The global optimal proportional routing problem (1) is a *constrained nonlinear optimization problem* and can be solved using an iterative procedure based on the *sequential quadratic programming method* [18], [3]. Each stage of the iterative procedure has two steps. First, given a set of flow proportions α_r , the fixed-point equations (3) and (4) involving b_l 's and ν_l 's are solved. Using these values, W given by

¹The models presented in this note can be extended to the case where flows have different bandwidth requirements using the extended Erlang loss formula [7], [19].

(1) is recomputed. Then this algorithm essentially searches for a new set of improved flow proportions based on the revenue W .

B. Virtual Capacity Model

We now turn our attention to the problem of modeling *localized* proportional routing. Unlike in global proportional routing, in localized proportional routing we assume that each source node has only a *local* (and thus *partial*) view of the network state. For example, a source node may only have knowledge of the offered traffic loads between the source–destination pairs originating from itself. It may also only have partial network topology information (in particular, the link capacity information may not be available to a source node). As mentioned in the introduction, in this note we will focus on local QoS state information gathered at two different granularity levels: the *link* level and the *path* level. At the (finer) link level, each source node can collect the following information locally: 1) the offered traffic load of flows from the source to a destination; 2) the number of flows routed along a path from the source to a destination that are blocked; and 3) in the case when a flow is blocked, the identity of the link at which the flow is blocked. The third type of information can be made available to a source node by simply piggybacking the identity of the link at which a flow is blocked in the flow setup failure notification sent back to the source node. At the (coarser) path level we assume that each source node only collects the first and second types of the local information listed above. As a result, the path-level local information provides a source node with a much “vaguer” view of the global network QoS state.

Given only locally collected flow statistics, determining “optimal” proportions for distributing flows among multiple paths between a source–destination pair becomes a difficult problem. In particular, since each source node does *not* know the capacity of a link and the total offered load on the link, the Erlang loss formula cannot be directly used to derive flow blocking probability at a link. To address this problem, we introduce the notion of *virtual capacity* of a link (or a path) *perceived* by a source node. Consider a link l , let $\nu_{s,l}$ be the load placed by a source node s , and $b_{s,l}$ be the blocking probability observed by node s . Intuitively, the virtual capacity, $vc_{s,l}$, of link l perceived by the source node s is the (perceived) amount of bandwidth available to the flows routed from source s along link l , given the observed blocking probability $b_{s,l}$. Formally, $vc_{s,l}$ is defined via the inverse of the Erlang loss formula as follows:

$$vc_{s,l} = E^{-1}(b_{s,l}, \nu_{s,l}) \quad (5)$$

where $E^{-1}(b, \nu) := \{c : E(\nu_{s,l}, c) = b_{s,l}\}$ is the inverse function of the Erlang loss formula with respect to the capacity. Note that we use the *continuous* version of the Erlang loss formula defined in [5]. The virtual capacity of a *path* can also be defined analogously by replacing the link l with a path r , $\nu_{l,s}$ and $b_{l,s}$ with $\nu_{r,s}$ and $b_{r,s}$, the load offered and blocking probability observed by node s along path r .

The notion of virtual capacity previously defined has several interesting and important properties that are key to our study of localized (adaptive) proportional routing. First of all, it is clear that the virtual capacity of a link or a path can be computed solely based on local information (e.g., load offered and blocking probability observed by a source node). Second, the notion of virtual capacity provides a *quantitative* measure of capacity² share on a link or a path grabbed by the flows originated from a source node. The larger the load a source node offers on a link or a path, the more capacity share the node grabs. We later see

²It is worth noting that $\sum_{i=1}^m vc_i \geq c$. This is due to “loss in multiplexing gain” when a shared channel is divided into multiple “dedicated” channels. To ensure the same blocking probability, the total capacity of the dedicated channels has to be larger than the capacity of the shared channel.

that this property helps the adaptation process by having a source node with fewer good candidate paths grab more capacity on a shared link, causing the other sources with better paths to adjust their proportions. Third, the virtual capacity perceived by a node is a function of *both* its offered load *and* the observed blocking probability, which changes as the *overall* load on a link or a path varies. Consequently, a node can adjust its offered load to effect a change in the observed blocking probability, or as a response to the change in the observed blocking probability. The notion of virtual capacity therefore provides a theoretical basis for the analysis of how flow proportions should be adjusted based on locally collected statistics.

III. LOCALIZED PROPORTIONAL ROUTING

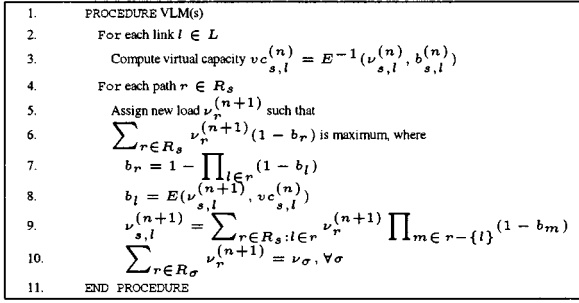
In this section, we present two virtual capacity based theoretical schemes for localized proportional routing—the *virtual link based minimization* (vlm) and the *virtual path based minimization* (vpm). In both schemes, each source collects local QoS state information, and based on this local QoS state information, periodically recomputes flow proportions assigned to the paths from the source to a destination. This distributed dynamic adaptation procedure can be viewed as an iterative process where in each iteration each source independently attempts to minimize the observed blocking probability by adjusting the amount of traffic routed through each path. These localized proportional routing schemes differ in the type and the granularity of local QoS state information collected, and therefore, in the computation of flow proportions for paths.

A. VLM

In the vlm model, a source collects link-level flow blocking statistics with the assistance from the connection admission control (CAC) module. We assume that whenever a flow setup request fails at a link, the identity of that link is also recorded and piggybacked to the source. The CAC module at the source node informs the QoS routing module of the flow setup failure and the identity of the link where the flow is blocked. Such link-level flow blocking information can be gathered by a source with very little overhead on the network.

With the locally collected link-level flow statistics, a source knows the offered traffic load on a link contributed by flows originating from that source. Unlike the global routing model, the source, however, does not have any information regarding the traffic loads offered by the other sources on the link. It neither has any knowledge of the capacity of the link. The source can only infer the state of the link from the flow blocking probability at the link it observes. Using the notion of virtual capacity of a link, the source can infer its share of the bandwidth at each link, and piece together a partial *virtual view* of the network from its own perspective.

With the virtual network view, each source can employ a localized version of the global optimal proportional routing scheme to compute the “optimal” flow proportions for each of its destinations: we replace the actual capacity of a link by its virtual link capacity, and only offered traffic loads from the source is used to compute the optimal flow proportions for the source. The resulting optimization procedure, referred to as the virtual link based minimization (vlm) procedure, is shown in Fig. 1, where s is a source node. This localized flow proportioning scheme is an iterative process where each iteration is performed after an observation interval by each source asynchronously. In the n th iteration, the current virtual capacity $vc_{s,l}^{(n)}$ of each link l with respect to s , is computed, based on the current offered load $\nu_{s,l}^{(n)}$ and the corresponding observed blocking probability $b_{s,l}^{(n)}$. The local minimization is then performed on the virtual network thus formed with each link l having the capacity $vc_{s,l}^{(n)}$.

Fig. 1. The vlm procedure at source node s .

B. VPM


In the vpm model, each source collects only path-level flow statistics: the number of flows routed along each path between the source to a destination, and the number of flows blocked along the path. Unlike the link-level localized QoS routing model, here we assume that the identify of the link at which a flow is blocked is *not* available to a source. This, for example, will be the case if the link identity at which a flow is blocked is not piggybacked to a source as part of flow setup failure notification.

With only locally collected path-level flow statistics, a source can not infer the QoS state of any individual link. A source can only obtain some knowledge about the “quality” of a path based on the traffic offered on the path and the corresponding observed flow blocking probability along the path. Similar to the virtual link based QoS routing model, in the virtual path based routing model we associate a virtual network with each source–destination pair, using the notion of virtual capacity of a path. Consider a source–destination pair (s, d) . Suppose there are k candidate paths between source s and destination d . Using the notion of virtual capacity of a path, we treat these k paths as if they were *disjoint* and each consisted of a single *virtual link*. The virtual capacity of a path r is represented by vc_r , which is determined by the offered load from source s to destination d along path r and the observed blocking probability b_r of flows routed along the path r . Although the real network topology of these paths may be very complex (e.g., multiple paths of a pair may share links among them or with paths of other source–destination pairs), the notion of virtual capacity of a path allows us to circumvent these difficulties by essentially capturing the “capacity share” of flows routed along various paths.

Given the path-level virtual network view for a source–destination pair, the “optimal” flow proportions for the paths between the pair can be computed to minimize the overall flow blocking probability experienced by the flows routed along these paths. Formally, consider a source–destination pair σ . Let R_σ denote the set of paths between the source–destination pair σ . For each path $r \in R_\sigma$ let vc_r denotes its virtual capacity (perceived by the source–destination pair σ). The flow proportions for the paths can be computed using an iterative procedure, referred to as the virtual path based minimization (vpm) procedure, as is shown in Fig. 2. In this procedure, the virtual capacity $vc_r^{(n)}$ of each path r is computed using the Erlang inverse formula, given the current offered load $\nu_r^{(n)}$ along the path r and the corresponding observed blocking probability $b_r^{(n)}$. Based on these path virtual capacities, new loads $\{\nu_r^{(n+1)}\}$ are reassigned to paths such that $\sum_{r \in R_\sigma} \nu_r^{(n+1)} E(vc_r^{(n)}, \nu_r^{(n+1)})$ is minimized. This procedure is performed iteratively and independently at each source per each destination.

C. Alternative Paths and Localized Trunk Reservation

The virtual capacity based local minimization schemes described so far treat all candidate paths equally. Since an admitted flow consumes

Fig. 2. The vpm procedure for a pair σ .

bandwidth and buffer resources at all the links along a path, clearly, path length is also an important factor that must be taken into consideration. There is a fundamental tradeoff between minimizing the resource usage by choosing shorter paths and balancing the network load by using lightly loaded longer paths. As a general principle, it is preferable to route a flow along *minhop* (i.e., shortest) paths than paths of longer length (also referred to as *alternative* paths). By preferring minhop paths and discriminating against alternative paths, we not only reduce the overall resource usage but also limit the so-called “knock-on” effect [8], [9], thereby ensuring the stability of the whole system. The knock-on effect refers to the phenomenon where using alternative paths by some sources forces other sources whose minhop paths share links with these alternative paths to also use alternative paths. This cascading effect can cause a drastic reduction in the overall throughput of the network.

In order to deal with the knock-on effect, trunk reservation [9] is employed where a certain amount of bandwidth on a link is reserved for minhop paths only. A flow along a path longer than its minhop path is admitted only if the available bandwidth even after admitting this flow is greater than the amount of trunk reserved. Trunk reservation provides a simple and yet effective mechanism to control the knock-on effect. However, trunk reservation cannot be used directly in localized routing schemes, since it requires global configuration. Furthermore, core routers have to figure out whether a setup request for a flow is sent along its minhop path or not, introducing undesirable burden on them. We propose to address this by having each source router locally discriminate against its own alternative paths *without any explicit global trunk reservation*. A source node employing a localized scheme can control the amount of alternative routing by adjusting the virtual capacities in its virtual network. This can be thought of as an *implicit localized trunk reservation* performed by each source independently.

The exact method in which alternative paths are discriminated varies between vlm and vpm. While vlm employs link-level discrimination, vpm does path-level discrimination. Under vlm with trunk reservation, a link is categorized into two cases: *alternative-only* or *minhop-also*. A link l is said to be alternative-only link w.r.t. a source s , if l lies only along alternative paths from the source s to its destinations. Otherwise if a minhop path from source s to any destination passes through link l , then l is categorized as minhop, also w.r.t. source s . The links that are used only by alternative paths for routing traffic from this source are targeted for the adjustment. Their capacities are reduced by an amount ψ where, ψ is the trunk reservation parameter, i.e., $vc_{s,l} = (1 - \psi)vc_{s,l}$ if l is alternative-only w.r.t. source s . The capacities of other links are left unchanged. In case of vpm, this local adjustment of virtual capacities is more straightforward. Given a trunk reservation parameter ψ , the target virtual capacity of alternative paths is reduced by an amount ψ , i.e., $vc_r = (1 - \psi)vc_r$ if r is an alternative path. The virtual capacities of minhop paths are left unchanged. In both vlm and vpm, the minimization procedure is then applied locally at the source on the virtual network with these adjusted capacities.

Before we proceed to evaluate the performance of these schemes, it is interesting to contrast the link based and path based localized proportional routing models with the global optimal proportional routing model in the way they handle the sharing of links among paths. While the global model is aware of how the links are shared by all the paths between any source to any destination, the localized link-level model is only aware of sharing of links among the paths from the same source. The localized path-level model is completely oblivious of any link sharing. However, this lack of knowledge about explicit sharing between paths is somewhat compensated by the notion of virtual capacity, which indirectly accounts for the effect of link sharing. Moreover, the localized models make up for the absence of such knowledge by employing an iterative process to compute flow proportions, in an attempt to approach the optimal flow proportioning. This iterative procedure can be thought of as continual refinement of the (partial) virtual network view of each source that eventually converges to an equilibrium state yielding near-optimal flow proportions.

D. Performance Evaluation

In this section, we demonstrate the convergence process of the localized proportional routing models, and compare their blocking performance with that of the global optimal proportional routing model through numerical investigation. Before we present the results, we first describe the evaluation setup.

The *minisp* topology used in our study is shown in Fig. 3 which is the core of an ISP backbone topology used in [12]. For simplicity, all the links are assumed to be bidirectional and of equal capacity C in each direction. For each source, a subset of nodes (shown in smaller font) are chosen as destination nodes. The flow dynamics of the network are modeled as follows (similar to the model used in [22]). Each flow is assumed to require one unit of bandwidth. Flows arrive at a source node according to a Poisson process with rate λ . The incoming traffic at a source is uniformly split among its destination nodes. The holding time of a flow is exponentially distributed with mean $1/\mu$. Following [22], the offered network load is given by $\rho = \lambda N \bar{h} / \mu LC$, where N is the number of source nodes, L is the number of links, and \bar{h} is the mean number of hops per flow, averaged across all source-destination pairs. The parameters used in our study are $C = 20$, $\mu = 1$, $N = 9$, $L = 26$, and $\bar{h} = 2.64$. The average arrival rate at a source node λ is set depending upon the desired load.

We first illustrate the convergence of localized schemes. The load ρ is set to 0.60 and only the minhop paths are chosen as the candidate paths for each source-destination pair. The average period between recomputations is set to 1. The Fig. 4(a) shows the overall blocking probability as a function of time, i.e., the number of iterations. The performance of the global optimal routing scheme is also shown for reference. Note that the performance of localized schemes only varies with time. It can be seen that the overall blocking probability of both the localized routing schemes gradually decreases as the number of iterations increases. Both the schemes eventually converge and each to a different convergence point. Starting with arbitrary initial proportions, the localized schemes approach close to their respective convergence points within ten iterations. Though the finer-grained link-level scheme performs better than the coarser-grained path-level scheme, there is not significant difference in their blocking probabilities. More importantly, the blocking performance of both the localized schemes is quite close to that of the global optimal scheme.

We now study the effectiveness of localized trunk reservation method and the impact of parameter ψ on the performance. Apart from *minhop* paths, paths of length $\text{minhop} + 1$ are also chosen as the candidate paths for each source-destination pair. Fig. 4(b) and (c) show the convergence process of vlm and vpm, respectively. The




Fig. 3. The *minisp* topology used in our study.

performance of these schemes is shown for different values of the trunk reservation parameter ψ : 0%, 5%, 10%, 15%. It is quite evident that the performance of localized schemes with ($>0\%$) trunk reservation is better than without (0%) it. However, as the ψ value is increased the performance gain is reduced. There is almost no difference in performance between ψ values of 10% and 15%. These results show that localized trunk reservation is quite effective, particularly in case of vpm. This is expected since vlm even without any trunk reservation, using finer-grained link-level information, accounts for sharing of links between minhop and alternative paths from a source to all its destinations. The role of localized trunk reservation in vlm is limited to avoiding overloading of minhop paths of a source by traffic on alternative paths of another source. On the other hand, the localized trunk reservation plays a much more critical role in vpm. The minhop paths to a destination have to be guarded from alternative paths to the same destination besides from alternative paths to other destinations. This is due to availability of only coarser-grain information, and thus, lack of knowledge about sharing of links between different paths. However, with localized trunk reservation, the vpm scheme tides over this shortcoming and performs comparably to the vlm scheme. We now proceed to study the practical issues involved in implementing localized schemes.

IV. ISSUES IN IMPLEMENTATION OF LOCALIZED SCHEMES

The localized schemes described so far have been shown to approach the performance of the optimal scheme using only local information. Furthermore, even with coarser-grain path-level blocking information the vpm scheme performs as well as the vlm scheme that uses finer-grain link-level blocking information. It is easier to collect path-level statistics and simpler to implement path-based schemes. Hence, we focus on path-based localized schemes and further investigate the issues involved in implementing them.

A. Equalization-Based Proportioning

The vpm scheme first computes the virtual capacity of each candidate path and then performs local minimization. Though the complexity of this minimization procedure is much less than that of optimal scheme, it could still be significant. A simple alternative to minimization procedure is to *equalize* either blocking probabilities or blocking rates. The objective of the *equalization of blocking probabilities* (ebp) strategy is to find a set of proportions $\{\hat{\alpha}_1, \hat{\alpha}_2, \dots, \hat{\alpha}_k\}$ such that flow blocking probabilities of all the paths are equalized, i.e., $\tilde{b}_1 = \tilde{b}_2 = \dots = \tilde{b}_k$, where \tilde{b}_i is the flow blocking probability of path r_i , and is given by $E(\hat{\alpha}_i \nu, c_i)$. On the other hand, the objective of the *equalization of blocking rates* (ebr) strategy is to find a set of proportions $\{\hat{\alpha}_1, \hat{\alpha}_2, \dots, \hat{\alpha}_k\}$ such that flow blocking rates of all the paths are equalized, i.e., $\hat{\alpha}_1 \tilde{b}_1 = \hat{\alpha}_2 \tilde{b}_2 = \dots = \hat{\alpha}_k \tilde{b}_k$, where \tilde{b}_i is the flow blocking probability of path r_i , and is given by $E(\hat{\alpha}_i \nu, c_i)$.

The virtual path based equalization (vpe) procedure for equalizing blocking rates is shown in Fig. 5. At any given iteration $n \geq 0$, let $\nu_r^{(n)}$ be the amount of the load currently routed along a path $r \in R_\sigma$, and let




Fig. 4. Convergence of localized schemes (load = 0.60). (a) Minhop paths only. (b) vlm with alternative paths. (c) vpm with alternative paths.

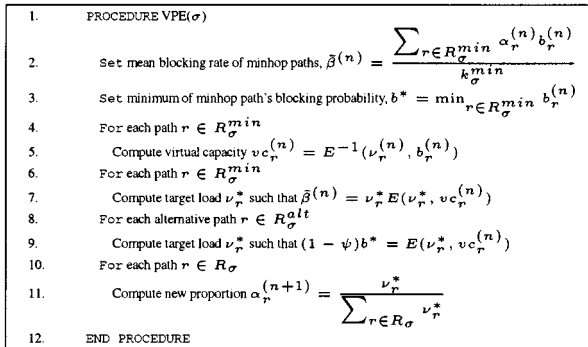


Fig. 5. The vpe procedure for a source–destination pair σ .

$b_r^{(n)}$ be its observed blocking probability on the path. Then the virtual capacity of path r is given by $v_{c_r} = E^{-1}(\nu_r^{(n)}, b_r^{(n)})$. For each minhop path, the mean blocking rate of all the minhop paths, $\bar{\beta}^{(n)}$, is used to compute a new target load. Similarly, for each alternative path, a new target load is computed using the target blocking probability $(1 - \psi)b^*$. Here, b^* is the minimum flow blocking probability of all the minhop paths and ψ is a configurable parameter to limit the knock-on effect. The basic idea behind this alternative routing method is to ensure that an alternative path is used to route flows between the source–destination pair only if it is better than all the minhop paths. Given these new target loads for all the paths, the new proportion of flows, $\alpha_r^{(n+1)}$, for each path r is obtained, resulting in a new load $\nu_r^{(n+1)} = \alpha_r^{(n+1)} \nu_{\sigma}$ on path r .

We now compare the performance of the vpe scheme with the vpm scheme. The Fig. 6(a) shows the performance of these schemes when only minhop paths are chosen as the candidate paths while the Fig. 6(b) shows the case when alternative paths are also used for routing. In both cases, the load is varied from 0.5 to 0.6 and the overall blocking probability is shown as a function of load. It can be seen that though the vpm performs better than the vpe scheme, the difference is not large. Moreover, with alternative paths there is almost no difference in the





Fig. 6. Performance of localized schemes under various loads. (a) Minhop paths only. (b) Alternative paths also.

performance of vpe and vpm. This can be attributed to the way the vpe scheme discriminates alternative paths using the performance of minhop paths as a reference.

The localized routing schemes presented so far are based on theoretical virtual capacity model. We have shown that they yield near-optimal performance using only local information. However, computation of virtual capacity and target load using Erlang's Loss Formula can be cumbersome. More importantly, the accuracy in using Erlang's Loss Formula to compute virtual capacity and new load relies critically on steady-state observation of flow blocking probability. Hence, small statistic variations may lead to erroneous flow proportioning, causing undesirable load fluctuations. To circumvent these difficulties, we have proposed simple yet robust implementations of these schemes, details of which can be found in [17].

B. Heterogeneous Traffic

The discussion so far is focused on the case where the traffic is homogeneous, i.e., all flows request for one unit of bandwidth. In [20], it was shown that when the capacity of a link is large, the blocking probability of a flow of type i can be approximated as follows. Suppose that type i flow requests for d_i units of bandwidth and the load of type i flows on link l is ν_i^l . The blocking probability for type i flows on link l is given by $b_i^l = (d_i/\delta)E((\sum \nu_i^l d_i)/(\delta), (c_l/\delta))$, where δ is an "equivalent rate" given by $\delta = (\sum \nu_i^l d_i^2)/(\sum \nu_i^l d_i)$. In other words, the ratio of blocking probabilities of flow types i and j would be same as the ratio of their bandwidth requests, i.e., $(b_i/b_j) \approx (d_i/d_j)$. This implies that $(\lambda_1 b_1)/(\lambda_2 b_2) = (\phi_1/\phi_2)$, i.e., the blocking rate of flows of a type is proportional to their fraction in the total offered load. Consequently, performance of an equalization based proportional routing scheme would be same with or without categorizing the flows into different classes. Considering that in practice link capacities are much larger than an individual flow's bandwidth request, proportional routing schemes can be used *as is* to route heterogeneous traffic.

C. Candidate Path Selection

The localized routing schemes discussed thus far are concerned only with computing proportions given a set of candidate paths. While the proportions for candidate paths are adjusted to reflect the changing network conditions, the candidate path set itself remains static. To reduce setup overhead and to ensure faster convergence, it is desirable to identify a few good candidate paths. However, due to changing network conditions, it is not possible to preselect a few good paths statically. Hence it is necessary to supplement localized proportional routing with a mechanism that dynamically selects a few good candidate paths. We have proposed such a hybrid approach to proportional routing where candidate paths are selected based on *infrequently* exchanged global link-level information and traffic is proportioned among candidate paths using locally collected path-level information. For more details, please see [16] and [17].

V. CONCLUSION

In this note, we set out to investigate an important and fundamental issue in localized QoS routing: the granularity of locally collected QoS state information and its impact on the design of localized QoS routing schemes. Toward this goal, we developed theoretical models for studying localized proportional QoS routing. These models are designed using the key notion of virtual capacity of a link or path, which provides the basis for a source to infer the perceived quality of a link or path based on only locally collected QoS state information.

Through numerical study, we demonstrated that although localized QoS routing schemes that make routing decisions based on locally collected QoS state information does pay a performance penalty for the partial and "vague" view of the global network QoS state, this penalty is not very significant. We have also discussed implementation issues and argued that by using simple localized equalization based proportioning strategies, it is possible to route heterogeneous traffic and yield near-optimal blocking performance with minimal communication overhead.

REFERENCES

- [1] R. Callon, A. Viswanathan, and E. Rosen, "Multiprotocol label switching architecture," *RFC 3031*, Jan. 2001.
- [2] S. Chen and K. Nahrstedt, "An overview of quality-of-service routing for the next generation high-speed networks: Problems and solutions," *IEEE Network Mag., Special Issue Trans. Distrib. Digital Video*, vol. 12, pp. 64–79, Nov./Dec. 1998.
- [3] T. Coleman and Y. Li, "An interior, trust region approach for nonlinear minimization subject to bounds," *SIAM J. Optim.*, vol. 6, pp. 418–445, 1996.
- [4] E. Crawley, R. Nair, B. Rajagopalan, and H. Sandick, "A framework for QoS-based routing in the Internet," *RFC 2386*, Aug. 1998.
- [5] R. F. Farmer and I. Kaufman, "On the numerical evaluation of some basic traffic formulae," in *Networks*. New York: Wiley, 1978, vol. 8, pp. 153–186.
- [6] R. Guerin, S. Kamat, A. Orda, T. Przygienda, and D. Williams, "QoS routing mechanisms and OSPF extensions," *RFC 2676*, Aug. 1999.
- [7] J. S. Kaufman, "Blocking in a shared resource environment," *IEEE Trans. Commun.*, vol. COM-29, pp. 1474–1481, 1981.
- [8] F. P. Kelly, "Routing in circuit-switched networks: Optimization, shadow prices and decentralization," *Adv. Appl. Probab.*, vol. 20, pp. 112–144, 1988.
- [9] —, "Routing and capacity allocation in networks with trunk reservation," *Math. Oper. Res.*, vol. 15, pp. 771–793, 1990.
- [10] —, "Fixed point models of loss networks," *J. Austr. Math. Soc.*, ser. B, vol. 31, pp. 204–218, 1989.
- [11] P. B. Key and G. A. Cope, "Distributed dynamic routing schemes," *IEEE Commun. Mag.*, vol. 28, pp. 54–64, Oct. 1990.
- [12] Q. Ma and P. Steenkiste, "On path selection for traffic with bandwidth guarantees," presented at the *IEEE Proc. ICNP*, 1997.
- [13] D. Mitra, J. A. Morrison, and K. G. Ramakrishnan, "ATM network design and optimization: A multirate loss network framework," *IEEE/ACM Trans. Networking*, vol. 4, pp. 531–543, Aug. 1996.
- [14] S. Nelakuditi, R. P. Tsang, and Z.-L. Zhang, "Quality-of-service routing without global information exchange," in *Proc. IWQOS*, 1999.
- [15] S. Nelakuditi, Z.-L. Zhang, and R. P. Tsang, "Adaptive proportional routing: A localized QoS routing approach," in *IEEE Proc. IN-FOCOM'00*, Mar. 2000.
- [16] S. Nelakuditi and Z.-L. Zhang, "On selection of paths for multipath routing," in *Proc. IWQOS'01*, June 2001.
- [17] S. Nelakuditi, "Localized Approach to Providing Quality-of-Service," Ph.D. dissertation, Univ. Minnesota, Minneapolis, MN, Oct. 2001.
- [18] M. Powell, "A fast algorithm for nonlinear constrained optimization calculations," in *Numerical Analysis, Lecture Notes in Mathematics*, G. A. Watson, Ed. New York: Springer-Verlag, 1978, vol. 630.
- [19] J. W. Roberts, "Teletraffic models for the telecom 1 integrated services network," presented at the *Proc. Internet Teletraffic Congress-10*. Session 1.1, paper #2.
- [20] J. Roberts, U. Mocci, and J. Virtamo, "Broadband network teletraffic," in *LNCS 1155*. New York: Springer-Verlag, 1996.
- [21] K. W. Ross, *Multiservice Loss Models for Broadband Telecommunication Networks*. New York: Springer-Verlag, 1995.
- [22] A. Shaikh, J. Rexford, and K. Shin, "Evaluating the overheads of source-directed quality-of-service routing," in *Proc. ICNP*, 1998.
- [23] Z. Zhang, C. Sanchez, B. Salkewicz, and E. Crawley, "Quality of service extensions to OSPF," unpublished, Sept. 1997.