Congestion control with adaptive multipath routing based on optimization

Fernando Paganini
ORT University,
Montevideo, Uruguay

***Abstract*— The paper considers a TCP/IP-style network with flow control at end-systems based on congestion feedback, and routing decisions at network nodes on a per-destination basis. The only generalization with respect to standard IP is that routers split their traffic, for each destination, among their outgoing links.**

We pose two optimization problems, that generalize and combine those used in the congestion control and traffic engineering literature. In contrast to other work in multipath congestion control, we use variables that are available at each node (source or router). We prove that decentralized algorithms built by combining primal or dual congestion control with adaptation of router splits, converge globally to optimal points. Some comments on practical implications of these results are given.

I. Introduction

The use of optimization and economic ideas in network resource allocation has a long history, going back to the study of transportation networks [21]. A classical problem is to minimize a cost of the form

$$\sum_l \phi_l(y_l) \qquad (1)$$

where $\{y_l\}$ represent flows in links of a network, subject to external traffic demands and flow conservation constraints at network nodes, and $\phi_l(\cdot)$ is a cost function, often taken to represent delay. A question that has been extensively studied (see [15] and references therein) are the equilibria that result from selfish routing decisions by traffic agents, and their inefficiencies in regard to the above social cost.

In IP networks, where routing decisions are made by routers, this type of optimization has been applied to problems of *traffic engineering* [6], [17], where a network operator seeks routes to serve a “traffic matrix” of demand. This is often done offline and subsequently implemented by some means (creation of MPLS tunnels or selection of weights in IP routing). In [5], an adaptive method based on optimization is proposed, running at the access points to an MPLS network, for a real-time optimization of (1). Most relevant to this paper is the work of Gallager [7], continued in [2], where network nodes perform an adaptive minimization of delay by controlling the split of traffic through outgoing links.

Email: paganini@ort.edu.uy. The author is on leave from the University of California, Los Angeles. Research supported in part by the David and Lucille Packard Foundation

A complementary line of research on *congestion control* has used optimization tools to control demand into the network. The basic such problem, proposed by Kelly [9], is the maximization of

$$\sum_k U_k(x^k) \qquad (2)$$

where x^k is the input rate of a traffic source, and $U_k(\cdot)$ an associated utility function. For single-path routing, there has been substantial progress in finding decentralized algorithms at sources and links to solve this problem, and relating these algorithms to current TCP congestion control (see [9], [12], [18], [13]).

Compared to this extensive research on either the supply or the demand sides of the problem, their combination (adapting both routing and source traffic) has been less studied. If single-path routing is imposed as in standard IP, it complicates the optimization of (2) through non-convexity [20]. However if we allow multiple routes, the problem is well-behaved, as was already noted in [9]. Here, and in [8], [19], [11], the proposal is to use as adaptation variables the components of rate for *each path* from source to destination. This gives convergent decentralized algorithms, but appears to be an impractical proposition in a large network. Achieving optimality would require sources to separately control rate on an exponential number of end-to-end paths, based on all combinations of routing choices along the way. This is not scalable, and constraining the set of paths will reduce utility. Also, propagating this control inside the network requires source routing, a significant departure from current Internet practice.

In this paper we propose to use only adaptation variables with local meaning: source rates, link congestion prices, and the traffic split at each router among its outgoing links for each destination. This combines congestion control with the adaptive routing of [7]. We will present natural optimization problems, and study how they can be solved in a decentralized way by primal and dual congestion control algorithms combined with a suitably chosen adaptation of traffic splits.

Other related work we recently became aware of is the upcoming paper [4], that formulates a very similar optimization problem for wireless scheduling. Given the differences in their dynamic solution and implementation context, we do not know at present how both compare.

II. Problem formulation

We consider a network made up of a set of nodes $\mathcal{N}$, and a set of directed links $\mathcal{L}$ between them. Nodes, denoted by the indices i and j, can be sources or destinations of packets, or intermediate router nodes. We describe the links either by a single index l, or by the directed pair (i, j) of nodes they connect.

The network supports various flows between source-destination pairs of nodes. We use the index $k \in \mathcal{K}$ to denote an individual flow or "commodity", and $s(k)$, $d(k)$ denote respectively the corresponding source and destination nodes. While these are unique for each k, we allow the traffic to follow multiple paths between source and destination. This is modeled through the following variables for each k: x^k, external flow in packets per second entering the network at the source; y_l^k, flow through link l; x_i^k, total flow coming into node i.

At the source node, we have

$$x_{s(k)}^k = x^k, \tag{3}$$

which assumes no commodity k traffic loops back to the source. We also write the flow balance equations

$$x_j^k = \sum_{(i,j)\in\mathcal{L}} y_{i,j}^k, \quad j \neq s(k), \tag{4}$$

$$x_i^k = \sum_{(i,j)\in\mathcal{L}} y_{i,j}^k, \quad i \neq d(k). \tag{5}$$

The total flow on link l is given by

$$y_l = \sum_k y_l^k. \tag{6}$$

Following [9], we will associate with each commodity k an increasing, concave utility function $U_k(x^k)$ that specifies the traffic's demand for rate. We formulate the following multipath counterpart of the "system problem" in [9].

Problem 1 (SYSTEM): Maximize $\sum_k U_k(x^k)$, subject to link capacity constraints $y_l \leq c_l$, and flow balance constraints (3),(4),(5),(6).

The solution of this convex program gives the maximum achievable utility over all sources if traffic is allowed to follow multiple routes between source and destination.

A second problem, also considered in the single-route case, can be formulated replacing capacity constraints with by barrier functions $\phi_l(y_l)$ that specify the congestion cost at the link. We assume $\phi_l(y_l)$ is increasing and convex in y_l. The barrier function problem is specified as follows.

Problem 2 (BARRIER): Maximize

$$S := \sum_k U_k(x^k) - \sum_l \phi_l(y_l) \tag{7}$$

subject to flow balance constraints (3),(4),(5),(6).

The above convex program combines the utility maximization of (2) with the cost minimization of (1). As such, it combines the congestion control formulation with tools used in traffic engineering. In economic terms, the quantity S defined in (7) is the *aggregate surplus* (see e.g. [14]), and a natural object of optimization; see also [16] for a similar object in a transportation context. To be meaningful, utility and cost must be expressed in the same units of "money" (a congestion currency relevant to all entities in the network). We will assume the cost functions grow fast enough for large y_l so that surplus is upper bounded, and Problem 2 has a finite optimum.

By appropriate redefinition of the variables, the above problems can be shown to be equivalent to those considered in [9], [11], in terms of rates per route or path. As argued before, we prefer to use variables which are fewer and directly accessible to either sources or routers.

In the Appendix, the solutions to these problems are characterized through Lagrangian duality.

III. Decentralized implementation

The challenge is to find decentralized solutions to these problems that can be embedded in network sources and routers. To respect as much as possible the information constraints of the Internet, we will work with:

- IP-like routers, which make routing decisions based on packet destination only. The main generalization we allow here is the use of multiple outgoing links with control over the traffic split.
- TCP-like sources that control transmitted rate based on a simple congestion feedback signal.

A. Control variables

First, following [7] we introduce at each node i traffic split variables for each destination d, satisfying

$$\alpha_{i,j}^d \geq 0, \qquad \sum_{(i,j)\in\mathcal{L}} \alpha_{i,j}^d = 1.$$

We will constrain our system dynamics by

$$y_{i,j}^k = \alpha_{i,j}^{d(k)} x_i^k, \qquad (i,j) \in \mathcal{L}; \tag{8}$$

in words, $\alpha_{i,j}^{d(k)}$ controls the fraction of incoming traffic of commodity k that node i sends through link $l = (i, j)$.

An alternative would be to follow the splits only on a per-destination basis, i.e. to impose

$$\sum_{d(k)=d} y_{i,j}^k = \alpha_{i,j}^d \sum_{d(k)=d} x_i^k, \qquad (i,j) \in \mathcal{L}; \tag{9}$$

This weaker restriction may be advantageous in practice, allowing nodes to route each flow as much as possible through a single route, as long as totals per destination are adequately split. While we will use (8) for simplicity in the presentation, the theory in this paper extends to flows satisfying (9); in particular, both alternatives lead to optimal allocations.

Next, we define a congestion measure or *price* p_l for each link $l \in \mathcal{L}$; we assume p_l depends only the total traffic y_l; there is no "service differentiation" between commodities.

A TCP-like source should not have to manage congestion information in different routes inside the network; rather, it should receive a simple signal that reflects the entire congestion state of all paths available to it. For this purpose, first define the node prices q_i^d, $i \in \mathcal{N}$, representing the average price of sending packets from node i to destination d, using the current routing patterns. Node prices are thus chosen to satisfy

$$q_d^d = 0,$$
$$q_i^d = \sum_{(i,j)\in\mathcal{L}} \alpha_{i,j}^d [p_{i,j} + q_j^d], \quad i \neq d. \tag{10}$$

Given link prices $p_{i,j}$, it follows through similar arguments as those in [7] that the above equations have unique solutions for q_i^d, provided that the split ratios α^d have a path from every node to the destination.

To determine the q_i^d in a decentralized network requires some communication across the nodes, and some time for prices to propagate. We do not model this process; further comments are given in Section V.

At the source node of commodity k, the node price summarizes the congestion cost of the network. We denote it by

$$q^k := q_{s(k)}^{d(k)}.$$

B. Basic relationships

The following basic lemma relates the price and flow variables defined so far.

Lemma 1: For each commodity k,

$$x^k q^k = \sum_{l\in\mathcal{L}} y_l^k p_l \tag{11}$$

Proof: We write the sequence of identities

$$\begin{aligned}
\sum_{i\in\mathcal{N}} x_i^k q_i^{d(k)} &= \sum_{i\in\mathcal{N}\setminus d(k)} x_i^k \sum_{(i,j)\in\mathcal{L}} \alpha_{i,j}^{d(k)} (p_{i,j} + q_j^{d(k)}) \\
&= \sum_{(i,j)\in\mathcal{L}} y_{i,j}^k (p_{i,j} + q_j^{d(k)}) \\
&= \sum_{l\in\mathcal{L}} y_l^k p_l + \sum_{j\in\mathcal{N}\setminus s(k)} q_j^{d(k)} \sum_{(i,j)\in\mathcal{L}} y_{i,j}^k \\
&= \sum_{l\in\mathcal{L}} y_l^k p_l + \sum_{j\in\mathcal{N}\setminus s(k)} q_j^{d(k)} x_j^k
\end{aligned}$$

The third step above follows by grouping terms by the end-nodes of the links. Now cancelling node terms, only the source term $x^k q^k$ remains on the left-hand side. ■

If we now aggregate the various commodities, we can use (11) and (6) to obtain the following corollary.

Corollary 2:

$$\sum_k x^k q^k = \sum_{l\in\mathcal{L}} y_l p_l \tag{12}$$

We now state dynamic relationships for the case when our variables (rates x, y, prices p, q, and split ratios α) vary in time; these hold regardless of the chosen control laws, to be defined later. The proof is omitted for brevity.

Lemma 3: For each commodity k,

$$\dot{x}^k q^k = \sum_{l\in\mathcal{L}} \dot{y}_l^k p_l - \sum_{(i,j)\in\mathcal{L}} x_i^k \dot{\alpha}_{i,j}^{d(k)} [p_{i,j} + q_j^{d(k)}] \tag{13}$$

$$x^k \dot{q}^k = \sum_{l\in\mathcal{L}} y_l^k \dot{p}_l + \sum_{(i,j)\in\mathcal{L}} x_i^k \dot{\alpha}_{i,j}^{d(k)} [p_{i,j} + q_j^{d(k)}] \tag{14}$$

$$\tag{15}$$

Corollary 4:

$$\sum_k \dot{x}^k q^k = \sum_{l\in\mathcal{L}} \dot{y}_l p_l - \sum_k \sum_{(i,j)\in\mathcal{L}} x_i^k \dot{\alpha}_{i,j}^{d(k)} [p_{i,j} + q_j^{d(k)}] \tag{16}$$

$$\sum_k x^k \dot{q}^k = \sum_{l\in\mathcal{L}} y_l \dot{p}_l + \sum_k \sum_{(i,j)\in\mathcal{L}} x_i^k \dot{\alpha}_{i,j}^{d(k)} [p_{i,j} + q_j^{d(k)}] \tag{17}$$

C. Route adaptation

We now discuss how to update $\alpha_{i,j}^d$. If the various prices that contribute to the average node price q_i^d are different, it is intuitive that the router should transfer traffic to cheaper routes; rather than shift *all* traffic to the cheapest path, which causes oscillations (see [20]), we gradually update the $\alpha_{i,j}^d$ in this direction, as in [7]. Specifically, for each destination d and node i the vector (over j) of derivatives $\{\dot{\alpha}_{i,j}^d\}$ should satisfy:

- The vector $\{\dot{\alpha}_{i,j}^d\}$ is a function of the vectors of current ratios $\{\alpha_{i,j}^d\}$ and the prices $\{p_{i,j} + q_j^d\}$.
- $\{\dot{\alpha}_{i,j}^d\}$ is negatively correlated with the route prices, and maintains node balance:

$$\sum_{(i,j)\in\mathcal{L}} \dot{\alpha}_{i,j}^d (p_{i,j} + q_j^d) \leq 0 \tag{18}$$

$$\sum_{(i,j)\in\mathcal{L}} \dot{\alpha}_{i,j}^d = 0 \tag{19}$$

- Equality in (18) occurs only if $\{\dot{\alpha}_{i,j}^d\} = 0$, and this happens only if for each $(i,j) \in \mathcal{L}$ we have

$$\begin{aligned}
&\text{either} \quad q_i^d = p_{i,j} + q_j^d, \\
&\quad \text{or} \quad \alpha_{i,j}^d = 0 \text{ and } q_i^d < p_{i,j} + q_j^d.
\end{aligned} \tag{20}$$

  In other words, split ratios per node only settle when prices of routes that carry traffic have equalized (and thus are equal to the node price) and the remaining unused routes have higher price.

We will not pick a specific route adaptation law; there is more than one choice that satisfy the above, all based on transferring traffic from more expensive to cheaper routes. See [7], [2] for some proposals in a discrete-time setting. Those papers also included a set of additional rules destined to keep the routing "loop free" at all times; we will not address this issue in this paper.

An important property of the above dynamics is that if link prices settle, so will split ratios and node prices. We omit the proof.

Proposition 5: Suppose $p_{i,j}$ are constant. Under the above assumptions for $\dot{\alpha}$, and q specified by (10), the variables α and q asymptotically reach equilibrium.

IV. Global convergence of congestion control with route adaptation

We now combine the route adaptation defined above, with two choices of congestion control algorithms studied in the literature. We use the notation

$$[w]_z^+ := \begin{cases} w, & \text{if } w > 0 \text{ or } z > 0; \\ 0 & \text{otherwise.} \end{cases}$$

A. Primal algorithm

Consider the scenario in which source rates are updated by the *primal* equations [9],

$$\dot{x}^k = \kappa(x^k)[U_k'(x^k) - q^k]_{x^k}^+ \tag{21}$$

where $\kappa(x^k) > 0$, and link prices follow the static law

$$p_l := \phi_l'(y_l), \tag{22}$$

i.e. the price is the marginal cost of the link.

The state variables of the system are here the source rates and node split ratios. We study the asymptotic behavior of the system.

Theorem 6: Under (21-22), and our assumptions on the adaptation of α, the system converges globally to a solution of Problem 2.

Proof: We take the derivative of the surplus along system trajectories,

$$\dot{S} = \sum_k U_k'(x^k)\dot{x}^k - \sum_l \phi_l'(y_l)\dot{y}_l \tag{23}$$

$$= \sum_k [U_k'(x^k) - q^k]\dot{x}^k + \sum_k q^k \dot{x}^k - \sum_l p_l \dot{y}_l \tag{24}$$

$$= \sum_k \kappa(x^k)[U_k'(x^k) - q^k][U_k'(x^k) - q^k]_{x^k}^+ - \sum_k \sum_{(i,j)\in\mathcal{L}} x_i^k \dot{\alpha}_{i,j}^{d(k)}(p_{i,j} + q_j^{d(k)}), \tag{25}$$

where we have invoked (21) and (16). Both of the above sums are non-negative, using (18); so we conclude that $\dot{S} \geq 0$, the surplus increases along trajectories.

Invoking the Lasalle invariance principle [10], the state trajectories will converge to an invariant set inside $\{(x,\alpha) : \dot{S} = 0\}$. Consider a trajectory satisfying $\dot{S} \equiv 0$. The first term in (25) implies that for each commodity

$$U_k'(x^k) = q^k, \text{or } x^k = 0 \text{ and } U_k'(x^k) < q^k. \tag{26}$$

In particular, $\dot{x}^k \equiv 0$ so the external rate is constant. Also, from (25) and the conditions (20), for each flow k and node $i \neq d(k)$, we have at any given time

$$\text{either} \quad x_i^k = 0, \quad \text{or} \quad \dot{\alpha}_{i,j}^{d(k)} = 0 \;\forall (i,j) \in \mathcal{L}; \tag{27}$$

i.e, either the node does not carry traffic of the given commodity, or the split ratios at that node are constant.

This means that while $\dot{S} = 0$, all *link* flows are constant and thus so are link prices $p_{i,j}$. Split ratios can continue to evolve at nodes that carry no destination traffic, and so will node prices q. It is indeed possible that a system trajectory could satisfy $\dot{S} = 0$ for an interval of time, and come out of this state later when the evolving node prices provide a cheaper, currently unused route. However, for a trajectory moving entirely within the set $\dot{S} = 0$, as stipulated in the Lasalle principle, link rates and prices must remain fixed forever. Invoking Proposition 5, node prices and split ratios converge asymptotically to an equilibrium satisfying the conditions in (20). Therefore, the only invariant points inside $\{\dot{S} = 0\}$ are those that satisfy both (26) and (20). We will show in the Appendix that a point satisfying these conditions is an optimum of Problem 2. ∎

B. Dual algorithm

We now consider the "dual" congestion control algorithm first proposed in [12], where link prices follow

$$\dot{p}_l = \gamma_l [y_l - c_l]_{p_l}^+. \tag{28}$$

Based on the received price q^k, the sources choose a rate $x^k = f_k(q^k)$ that instantaneously maximizes $U_k(x^k) - q^k x^k$, i.e. the chosen rate satisfies (26).

For the case of single routes per flow, the above is a gradient algorithm for the Lagrangian dual of the system problem, so rates converge globally to the optimum, while prices converge to the corresponding Lagrange multipliers. It will be convenient to state a slight generalization of this: a system problem in which we allow multiple paths but the split ratios are fixed.

Problem 3 (SYS-α): Given a set of split ratios $\alpha_{i,j}^d$, maximize $\sum_k U_k(x^k)$, subject to $y_l \leq c_l$, flow balance constraints (3),(4),(6), and splitting constraints (8).

This is also a convex problem, let us denote the maximum by $\Psi(\alpha)$. Problem 1 corresponds to computing $\max_\alpha \Psi(\alpha)$, which is not necessarily convex over α.

Focusing on Problem 3, write the Lagrangian with respect to the capacity constraints,

$$L(\alpha, p, x) = \sum_k U_k(x^k) + \sum_l p_l(c_l - y_l) \tag{29}$$

$$= \sum_k [U_k(x^k) - q^k x^k] + \sum_l p_l c_l \tag{30}$$

Here we have invoked Lemma 1. We now maximize over x for fixed α, p, defining

$$W(\alpha, p) := \max_x L(\alpha, p, x). \tag{31}$$

Note from (30) that the maximization over x is unconstrained, the prices q^k reflect all the flow balance constraints for fixed split ratios. By duality, we have

$$\Psi(\alpha) = \min_p W(\alpha, p) = \min_p \max_x L(\alpha, p, x). \tag{32}$$

Proposition 7: For fixed α, the dual algorithm solves Problem 3.

Proof: We compute the derivative of $W(\alpha, p)$ over trajectories of the dual; since x is instantaneously maximizing in (31), the Envelope Theorem (see [14]) gives

$$\dot{W} = \frac{\partial L}{\partial p}\dot{p} = -\sum_k \dot{q}^k x^k + \sum_l \dot{p}_l c_l.$$

We now invoke (17) with $\dot{\alpha}^d_{i,j} = 0$, and obtain

$$\dot{W} = \sum_l \dot{p}_l(-y_l + c_l) = -\sum_l \gamma_l [y_l - c_l]^+_{p_l}(y_l - c_l) \leq 0.$$

Thus $W(\alpha, p)$, decreases over trajectories to a point where every link satisfies

$$y_l = c_l, \text{ or } y_l < c_l \text{ and } p_l = 0. \tag{33}$$

This means link prices, together with the corresponding rates, are at a saddle point of (32); from duality, the equilibrium rates solve Problem 3.

■

Key to the above argument was the assumption that α remains constant; what happens now if the split ratios are updated as specified in Section III-C? Through the additional term in (17), plus the assumption (18), we find this contributes a *positive* term to $\dot{W}$, so its dynamic behavior over time is inconclusive. This may appear disappointing, but it must be the case: a sensible adaptation of the split ratios tries to *maximize* $\Psi(\alpha)$, to solve Problem 1. Its contribution to the Lagrangian is therefore in the increasing direction.

Consequently, if we adapt routes at the same time-scale as prices, the analysis of the Lagrangian does not allow us to predict the overall behavior. If, instead, we assume route adaptation is much slower so that the optimization over p can be seen as instantaneous, the system converges to the desired solution.

Theorem 8: For each set of split ratios α, define $\Psi(\alpha)$ by (32), and assume prices and rates take instantaneously their saddle point values. Updating α through (18-19), $\Psi(\alpha)$ converges to its global maximum, the solution to Problem 1.

Proof: We compute the derivative $\dot{\Psi}$ over trajectories satisfying (18-19), where for the current $\alpha(t)$ prices instantaneously minimize $W(\alpha, p)$, and rates instantaneously maximize $L(\alpha, p, x)$ for the given split ratios and prices. The Envelope Theorem implies that we can take derivatives directly on the Lagrangian for fixed prices and rates; using (30) and (17), we have

$$\begin{aligned}\dot{\Psi} &= -\sum_k x^k \dot{q}^k \\ &= -\sum_k \sum_{(i,j)\in\mathcal{L}} x^k_i \dot{\alpha}^{d(k)}_{i,j}[p_{i,j} + q^{d(k)}_j] \geq 0.\end{aligned}$$

Again invoking Lasalle's principle, trajectories converge to an invariant set of states inside $\{\alpha : \dot{\Psi} = 0\}$.

A similar reasoning as in Theorem 6 allows us to conclude that asymptotically we reach a point where (20) holds, in addition to (26) and (33). The methods of the Appendix imply such points are optima for Problem 1.

■

V. Practical implementation issues

We give a few comments on the applicability of this theory for resource allocation in a TCP/IP-like network.

A first point is the formation of node prices from link prices, based on equation (10). If they are to be determined *explicitly* by nodes that only interact with their neighbors, this requires an iteration. Here, each node periodically updates q^d_i to the right-hand side of (10), based on announcements from downstream links, and announces its new price to upstream neighbors. This extends routing announcements in IP to carry price information. Under the assumption of continued connectivity, it can be shown that this recursion converges.

Now, we have not modeled the dynamics of this price propagation, and assumed it to be instantaneous, which means the other variables (source rates, link prices, split ratios) evolve at a slower time-scale. Regarding sources, it does not appear practical to assume a TCP-like congestion control evolves more slowly than IP routing updates; the theory would thus only be applicable under the interpretation that utility functions represent the long-term demand of aggregates of sources, rather than individual TCP flows. This issue was already brought up in [7] where demands are considered slowly varying.

Nevertheless, there is the possibility of forming the source price *implicitly* through ECN packet marking, as is often considered in congestion control. Suppose, for this purpose, that each link marks packets with probability p_l, and these probabilities are small. Denote by $\tilde{q}^d_i$ the probability that it is marked on route from i to d. Since the traffic at the node is split with ratios $\alpha^d_{i,j}$, to first order we will have the recursion

$$\tilde{q}^d_i = \sum_{(i,j)\in\mathcal{L}} \alpha^d_{i,j}[p_{i,j} + \tilde{q}^d_j], \quad i \neq d. \tag{34}$$

exactly the same as (10). Therefore, ECN marking automatically finds the correct node prices, and sources will see (to first order) a marking probability that reflects its price q^k, analogously to standard congestion control.

This avoids the explicit generation of node prices to control the sources, but node prices q^d_i will still be needed for route adaptation. In the absence of ECN feedback from destinations to routers (which seems impractical, see [3] for a related proposal with probing for delay), we must default back to IP-style routing updates. Fortunately, since the time-scale of route adaptation can be slower, we can allow enough time for the convergence of the node pricing iteration.

The overall dynamic picture would then look like this:

- A fast time-scale for congestion control, that adapts source rates and link prices for fixed split ratios. Here ECN marking can provide the feedback signal. In essence, this solves Problem 3, or for the primal case a barrier problem with fixed split ratios.
- An intermediate time-scale for propagating congestion prices to nodes, based on IP-routing updates.
- A slow, traffic engineering time-scale for adaptation of split ratios based on the stabilized node prices.

VI. Conclusions and Future work

We have proposed optimization problems that maximize aggregate utility or surplus of a network under multipath routing. We have shown how a combination of a dynamic adaptation of split ratios at nodes, together with primal or dual congestion control laws at sources and links, converge to the global optima, under certain time-scale assumptions. From a practical perspective, the main advantage is that each entity (source, router) controls variables which are locally meaningful; the required feedback information can be implemented with the same philosophy of current protocols, an issue to be further studied in the future.

As mentioned in the introduction, recent work [4] in the context of wireless scheduling has posed an optimization problem similar to our Problem 1, but with various differences. A detailed comparison is open for future research, as well as the possible impact of our formulation in the wireless ad-hoc context.

Another interesting question is what happens when network routers, rather than seek social welfare, attempt to maximize some profit; these issues were studied in [1] for simple networks using path variables; we intend to pursue them under the framework of this paper.

Appendix: global optima

A characterization of optimal points via Lagrangian duality implies that each of our algorithms (dual for Problem 1, primal for Problem 2) asymptotically reach optimality. For reasons of space, we cover only Problem 2, rewriting it as follows:

$$\text{Maximize } S(x,y) = \sum_k U_k(x^k) - \sum_l \phi_l \left(\sum_k y_l^k \right) \tag{35}$$

in the variables $x = \{x^k\}_{k\in\mathcal{K}}$, and $y = \{y_l^k\}_{k\in\mathcal{K},l\in\mathcal{L}}$, subject to the constraints for every k, and $i \neq d(k)$:

$$\begin{aligned} g_i^k(x,y) &:= \sum_{(i,j)\in\mathcal{L}} y_{i,j}^k - x^k = 0, \quad i = s(k), \\ g_i^k(x,y) &:= \sum_{(i,j)\in\mathcal{L}} y_{i,j}^k - \sum_{(j,i)\in\mathcal{L}} y_{j,i}^k = 0, \quad i \neq s(k), \end{aligned} \tag{36}$$

We characterize the optimum using the dual

$$\min_{\lambda} \max_{x,y\geq 0} L^{BARR}(x,y,\lambda), \quad \text{where}$$

$$L^{BARR}(x,y,\lambda) = S(x,y) + \sum_k \sum_{i\neq d(k)} \lambda_i^k g_i^k(x,y).$$

The variables λ, x, y are at a saddle point of the dual if they satisfy (36) and for each k we have:

$$\frac{\partial L^{BARR}}{\partial x^k} = U_k'(x^k) - \lambda_{s(k)}^k = 0, \quad (\text{or} < 0 \text{ and } x^k = 0) \tag{37}$$

$$\frac{\partial L^{BARR}}{\partial y_{i,j}^k} = -\phi_l'(y_l) + \lambda_i^k - \lambda_j^k = 0, \quad (\text{or} < 0,\ y_{i,j}^k = 0), \quad j \neq d(k); \tag{38}$$

$$\frac{\partial L^{BARR}}{\partial y_{i,j}^k} = -\phi_l'(y_l) + \lambda_i^k = 0, \quad (\text{or} < 0 \text{ and } y_{i,j}^k = 0), \quad j = d(k). \tag{39}$$

Now refer back to Section IV-A; we showed that the system converges to a point where (26) and (20) are satisfied. Defining $\lambda_i^k = q_i^{d(k)}$, (26) implies (37). Also, recalling that $p_l = \phi_l'$, (20) implies (38-39). Therefore we are asymptotically at an optimum of Problem 2.

References

[1] D. Acemoglu and A. Ozdaglar, “Price Competition in Communication Networks”, to appear in IEEE INFOCOM 2006.

[2] D. Bertsekas, E. Gafni and R. Gallager, “Second derivative algorithms for minimum delay distributed routing in networks”, *IEEE Trans. on Comm.*, Vol Com-32 (8), pp. 911-919, 1984.

[3] V.S. Borkar and P.R. Kumar, “Dynamic Cesaro-Wardrop equilibration in networks”, *IEEE Trans. on Automatic Control*, Volume 48(3), pp. 382- 396, Mar 2003.

[4] L. Chen, S. H. Low, M. Chiang and J.C. Doyle, “Cross-layer congestion control, routing and scheduling design in ad-hoc wireless networks”, to appear in IEEE INFOCOM 2006.

[5] A. Elwalid, C. Jin, S. Low, and I. Widjaja, “MATE: MPLS Adaptive Traffic Engineering”, IEEE INFOCOM 2001.

[6] B. Fortz and M. Thorup,“Internet Traffic Engineering by Optimizing OSPF Weights”, IEEE INFOCOM 2000.

[7] R. G. Gallager, “A minimum delay routing algorithm using distributed computation”, *IEEE Trans. on Communicactions*, Vol Com-25 (1), pp. 73-85, 1977.

[8] H. Han, S. Shakkottai, C.V. Hollot, R. Srikant and D. Towsley, Overlay TCP for Multi-Path Routing and Congestion Control. *ENS-INRIA ARC-TCP Workshop, Paris, France, Nov. 2003.*

[9] F. P. Kelly, A. Maulloo, and D. Tan, “Rate control for communication networks: Shadow prices, proportional fairness and stability”, *Jour. Oper. Res. Society*, vol. 49(3), pp 237-252, 1998.

[10] H. K. Khalil, *Nonlinear systems*, Prentice-Hall, 1996.

[11] I. Lestas and G. Vinnicombe “Combined control of routing and flow: a multipath routing approach”, *Proc. 43rd IEEE Conference on Decision and Control*, December 2004

[12] S. H. Low and D. E. Lapsley, “Optimization flow control, I: basic algorithm and convergence”, *IEEE/ACM Transactions on Networking*, vol.7, no.6,pp861-874, December 1999.

[13] S. Low, F. Paganini, J. Doyle, “Internet Congestion Control”, *IEEE Control Systems Magazine*, February 2002.

[14] A. Mas-Colell, M. Whinston, J. Green, *Microeconomic Theory*, Oxford University Press, 1995.

[15] T. Roughgarden, *Selfish Routing and the Price of Anarchy*, MIT Press, 2005.

[16] W. Sandholm, “Evolutionary Implementation and Congestion Pricing.” *Review of Economic Studies* 69 (2002), 667-689.

[17] A. Sridharan, R.Guerin, C. Diot. ”Achieving Near-Optimal Traffic Engineering Solutions for Current OSPF/ISIS Networks”. IEEE/ACM Transactions on Networking. March 2005

[18] R. Srikant, *The Mathematics of Internet Congestion Control*, Birkhauser, 2004.

[19] T. Voice, “A global stability result for primal-dual congestion control algorithms with routing”, *ACM Sigcomm CCR, 2004.*

[20] J. Wang, L. Li, S. H. Low and J. C. Doyle, “Cross-layer Optimization in TCP/IP Networks” *IEEE/ACM Trans. on Networking*, 13(3):582-268, June 2005.

[21] J. G. Wardrop, “Some theoretical aspects of road traffic research”, *Proc. Inst. of Civil Engineers*, pt II., vol 1, 325-378, 1952.