

# Trade & Cap

## *A Customer-Managed, Market-Based System for Trading Bandwidth Allowances at a Shared Link*

Jorge Londoño  
Computer Science Dept  
Boston University, USA  
jmlon@cs.bu.edu

Azer Bestavros  
Computer Science Dept  
Boston University, USA  
best@cs.bu.edu

Nikolaos Laoutaris  
Telefonica Research  
Barcelona, Spain  
nikos@tid.es

### ABSTRACT

We propose Trade & Cap (T&C), an economics-inspired mechanism that incentivizes users to voluntarily coordinate their consumption of the bandwidth of a shared network link so as to converge on what *they* perceive to be an equitable allocation, while ensuring efficient resource utilization. Under T&C, rather than acting as an arbiter, a service provider acts as an enforcer of what the community of rational users sharing the resource decides is a fair allocation of that resource. Our T&C mechanism proceeds in two phases. In the first, software agents acting on behalf of users engage in a strategic *trading* game in which each agent selfishly chooses *reserved* bandwidth slots to acquire in support of primary network usage activities. In the second phase, these agents acquire additional bandwidth slots in support of a presumed open-ended need for *fluid* bandwidth, catering to secondary applications. The acquisition of this fluid bandwidth is subject to the remaining “buying power” of each user and by prevalent “market prices” – both of which are determined by the outcomes of the trading phase, and by a desirable aggregate *cap* on link utilization. We present analytical results that establish the underpinnings of our T&C mechanism, including game-theoretic results pertaining to the trading phase, and pricing of fluid bandwidth allocation pertaining to the capping phase. Using Internet traffic traces, our experimental results demonstrate the benefits of our scheme, which we also show to be practical by highlighting the salient features of an efficient implementation architecture.

### 1. INTRODUCTION

**Motivation:** The ever increasing appetite for Peer-to-Peer (P2P), media streaming, and Video on Demand (VoD) content is forcing service providers to constantly upgrade their infrastructures to keep-up with customer bandwidth demands. This state-of-affairs is significantly exacerbated by the prevalence of flat-pricing schemes and hence the lack of an incen-

tive for users to moderate their hunger for network bandwidth, especially around periods of peak network utilization, which are the primary determinants of an Internet Service Provider (ISP) costs (both in terms of infrastructure upgrade cycle and inter-AS traffic volume costs due to the 95/5 rule). Attempts by ISPs to deviate from flat pricing (including field-tested per-byte pricing [2]) have been widely rejected by customers [13].

Under flat pricing, during periods of peak demand, current congestion control practices could be seen as particularly “unfair” to users of low-volume, mostly-interactive applications who are effectively subsidizing “bandwidth hogs.” This has prompted some ISPs to act as arbiters, proactively shaping user traffic by setting quotas,<sup>1</sup> or by preferentially treating different traffic payloads (*e.g.*, web browsing vs. bit-torrent downloads) during periods of peak demand.<sup>2</sup> These efforts have backfired, eliciting a public relation’s quagmire regarding violation of *Net Neutrality* [6].

**Scope and Contributions:** Rather than having ISPs act as arbiters who set *policies* regarding what constitutes fair usage of a shared resource, in this paper, we propose a *voluntary*, market-based Trade & Cap system in which user software agents converge to an allocation of resources that is perceived to be equitable by the community of users, irrespective of what these resources are used to support (HTTP vs P2P traffic) and irrespective of the absolute resource allocation (traffic volume) per user. In our setting, the role of the ISP is that of providing a *mechanism*, that supports any privately-defined user *policy* [5].

Effectively, our T&C mechanism sets up a marketplace. Given the fixed (*e.g.*, flat-rate) payment to the provider, customers enter this marketplace with equal buying power, but their use of this fairly-allocated buying power depends on their

---

<sup>1</sup>Incidentally, when demand is well below the provider’s nominal capacity, supporting bandwidth hogs is basically free, bringing to question the use of traffic “quotas” [19].

<sup>2</sup>There is a growing body of academic [1, 16, 26, 29] and industry [28, 9, 15] work on delineating different traffic payloads in order to police/balance consumption. Many of these systems depend on Deep Packet Inspection, raising privacy concerns. Moreover, the effectiveness of these techniques is questionable as applications adapt quickly to avoid detection, *e.g.* using encryption and port number randomization.

flexibility. This allows customers to trade “volume” during low-utilization periods for “quality” during peak-utilization periods (or vice versa). The direction of the trade (not to mention the willingness to even engage in trading) depends entirely on customer preferences and flexibility (*e.g.*, tolerance for delaying a scheduled network backup job).<sup>3</sup> In addition to empowering customers to trade bandwidth allocations, T&C has the desirable side effect of smoothing traffic utilization over time, thus reducing the ISP’s cost which is determined primarily by the peak rate.

**Outline and Summary of Results:** We start this paper in Section 2 by over-viewing the T&C mechanism as it applies to a Digital Subscriber Line Access Multiplexer (DSLAM) setting, and in Sections 3 and 4 by presenting analytical results pertaining to convergence and efficiency of the marketplace underlying T&C. In Section 5 we discuss the salient features of an implementation architecture for T&C in a DSLAM setting. Our implementation allows the marketplace interactions to be carried out by software agents that run on behalf of the users and of the ISP. With the exception of minimal configuration and parametrization, the actions of these agents is transparent to the user. Next, in Section 6, we demonstrate the significant advantages of T&C by presenting results from extensive trace-driven simulations. For instance, we show that introducing a relatively small level of flexibility in the scheduling of reserved bandwidth slots results in significant gains for both the users and the ISP. For example, allowing user agents to reposition a reserved bandwidth allocation within relatively small windows of time enables them to increase their share of fluid bandwidth allocation (supporting non-interactive applications) by 20% to 40% depending on their flexibility. This benefits the ISP as well, resulting in as much as 16% to 31% reduction in the 95th percentile of the ISP’s 5-minute traffic volume, and (even more impressively) resulting in smoothing traffic volume, reducing the 95th-percentile/50th-percentile ratio from 1.58 to an almost perfect ratio of 1.004. We conclude the paper in Section 7 with a review of the related literature.


## 2. TRADE & CAP IN A DSLAM SETTING

While our T&C mechanism is applicable to any setting in which it is desirable to coordinate the fractional acquisition by a set of self-interested parties of the *shared* capacity of a single resource, in this paper, and without loss of generality, we restrict ourselves to a specific setting – that of coordinating the utilization of a shared DSLAM uplink.

Figure 1 shows the basic architecture of Digital Subscriber Line (DSL) access technology. DSL modems on the customer side connect hundreds to thousands of users to a single DSLAM server on the provider network. DSLAMs connect to a Broadband Remote Access Server (BRAS) which relays traffic to/from the Internet. In this setting, a link in the path from the DSLAM to the upper tier provider poses significant traffic management problems for ISPs and is thus the shared resource managed using our T&C mechanism.<sup>4</sup>

<sup>3</sup>We take the liberty of using the terms *user(s)* and *customer(s)* to refer to the *customer-side software agent(s)* that engage in the T&C marketplace on behalf of the user(s).

<sup>4</sup>T&C is equally valuable and practical if the resource to be managed is not “physical” but rather “virtual” – *e.g.*, the aggregate inter-ISP (transit) traffic of a subnetwork. Our



**Figure 1: The DSL “last-mile” architecture and the T&C-managed DSLAM uplink.**

We assume that the marketplace operates over fixed, non-overlapping periods of time, which we call *epochs* (*e.g.*, days), and that the trading and allocation of capacity will occur within  $T$  subdivisions of an epoch, which we call *time slots*, *e.g.*, 288 5-minute slots per day to match a *de facto* industry standard of 5 minutes for traffic accounting and pricing.

At the start of each epoch, the operator assigns each agent  $i = 1, 2, \dots, n$  an allowance or *budget*  $B_i$  in accordance with the user’s Service Level Agreement (SLA) (*e.g.*, “Business” versus “Residential”). Under the flat pricing assumed in this paper, all customers receive an equal budget. Our T&C mechanism proceeds in two phases:

(1) *The Bandwidth Trading Phase:* This phase proceeds as a pure-strategies, non-cooperative game among agents, who are allowed to rationally and selfishly decide *when* to schedule bandwidth allocations in support of their Reserved Traffic (RT). Reserved Traffic is the traffic belonging to applications requiring a specific (minimum) bandwidth over a contiguous period of time. RT may be flexible in terms of start and end times but not in terms of the reserved bandwidth over time.<sup>5</sup> If RT is flexible, the agent’s goal would be to minimize the cost incurred in acquiring the fraction of the link’s capacity necessary to support the RT bandwidth. The scheduling of RT traffic is subject to preset user preferences and constraints. The outcome of this game is a Nash-Equilibrium (NE) of RT bandwidth allocations to all participating agents, along with the corresponding cost incurred by each agent.

(2) *The Bandwidth Capping Phase:* This phase proceeds as a market-clearing phase, in which the operator distributes any remaining capacity among agents. The amount of “remaining” capacity distributed in this phase is set based on a desirable nominal utilization of the link (*e.g.*, determined by the 95/5 rule threshold). The allocation of bandwidth in the capping phase rewards agents who were able to preserve more of their budgets in the trading phase (due to a low RT volume or due to flexibility in scheduling such traffic), ensuring a market equilibrium of the resulting allocations. We refer to the traffic due to applications that do not require specific reservations, but instead benefit by obtaining as much bandwidth as possible as Fluid-Traffic (FT) applications.

distributed implementation architecture discussed in Section 5 is particularly suited for managing such resources.

<sup>5</sup>The determination by an agent of what constitutes an appropriate reservation and associated flexibility could be based on local user traffic profiling – an orthogonal subject that has been studied widely, *e.g.* [23, 8, 27].

### 3. THE BANDWIDTH TRADING PHASE

Each agent  $i$  represents its RT demand as a vector of requested allocations:  $T_i = (t_{i1}, \dots, t_{iT_i})$ . An assignment of an agent's demand is a mapping that pins each one of the components of the vector to a different time slot. A set of such assignments (one per agent) comprises a potential configuration, or *schedule* of RT allocations at the DSLAM.

Let  $k = m_i(j)$  be the time slot assigned to the  $j^{\text{th}}$  component of agent  $i$ 's request vector. We denote by  $x_{ik}$  the actual allocation for agent  $i$  in time slot  $k$ , where  $x_{ik} = t_{i, m_i(j)}$ . The  $x_{ik}$  notation implicitly represents the mapping  $m_i(\cdot)$ , noting that for time slots that are not used in the mapping,  $x_{ik} = 0$ . Thus,  $x_{ik}$  is defined for all time slots.

DEFINITION 1. *The cost of the RT vector  $T_i$  is*

$$c_i = \frac{1}{C} \sum_{p=1}^T x_{ip} U_p \quad (1)$$

where  $U_p = \sum_{i=1}^n x_{ip}$  is the aggregate reservation on slot  $p$ , and  $C$  is a constant.

The motivation for the cost function in Equation 1 is two-fold. First, in schemes where cost is constant or proportional to the user's demand, there is no incentive for an agent to avoid congested time-slots – a given level of resource (bandwidth) usage costs the same in either case. Our cost function creates the desired incentive of steering agents away from congested time slots (if they possess the flexibility to do so). Second, our cost function is fair in the sense that users sharing the same time slot pay the same unit-price.

In the T&C marketplace, the cost  $c_i$  is deducted from the allowance, or budget  $B_i$  of agent  $i$ . The budget  $B_i$  reflects the "rights" allotted to agent  $i$  (*e.g.*, based on the user's service plan). Under a flat-rate service model, all users would be allotted equal budgets which they are free to use to acquire their RT and FT allocations.

The strategy space for agent  $i$  is the set of permutations of its request vector. As such, the strategy space is finite with cardinality  $P_i^T$ . The game's strategy space is the Cartesian product of the strategy spaces of all agents.

Notice that an agent may be subject to additional constraints that limit its strategy space – *e.g.*, a 2-hour-long RT fixed bandwidth allocation must be assigned in consecutive time-slots, and be scheduled to start between 6pm and 8pm. Two practical examples of such constraints are: (1) *Capacity constraints* to ensure that the shared link capacity is never exceeded by the aggregate allocation –  $\forall p : \sum_{i=1}^n x_{ip} \leq C$ , and (2) *Budget constraints* to ensure that no agent is able to reserve resources beyond its "fair" share, which is upper-bounded by the agent's allowance –  $\forall i : \frac{1}{C} \sum_{p=1}^T x_{ip} U_p \leq B_i$ .

THEOREM 1. *The pure strategies game in which agents adopt better/best responses to allocate their Reserved Traffic vectors converges to a Nash-Equilibrium.*

PROOF. Due to space limitations we refer the reader to the expanded version of this paper [21] for the proof.  $\square$

While instrumental in establishing the convergence property given in Theorem 1, the specification of a quadratic (square) form in our cost function in Equation 1 is not essential as other cost functions may well yield the same desirable incentive for agents to shift (if possible) their RT traffic allocation to lower-utilization time slots.<sup>6</sup>

### 4. THE BANDWIDTH CAPPING PHASE

The Capping Phase computes a market-clearing solution that allocates the left-over budget of the agents in such a way that maximizes the aggregate FT allocation for each user. Let  $w_{ip} \in \mathbb{R}^+$  be the allocation of FT for agent  $i$  in time-slot  $p$ . We adjust the definition of the cost function to take into account the allocation of FT as follows:

DEFINITION 2. *The cost to agent  $i$  for the combined allocation of RT ( $x_{ip}$ ) and FT ( $w_{ip}$ ) is*

$$c_i = \frac{1}{C} \sum_{p=1}^T (x_{ip} + w_{ip}) U_p \quad (2)$$

where  $U_p = \sum_{i=1}^n (x_{ip} + w_{ip})$  is the aggregate reservation on slot  $p$ , and  $C$  is a constant.

The implicit assumption of the Capping Phase is that RT allocations have priority, and are fixed once they are set at the conclusion of the the Trading Phase. FT allocations have no scheduling constraints: the value accrued by FT applications is strictly increasing with the aggregate allocation of FT bandwidth. Thus, self-interested agents select allocations so as to:

Maximize

$$\sum_{p=1}^T w_{ip}$$

subject to

$$c_i \leq B_i \quad (3)$$

$$w_{ip} \geq 0 \quad \text{for } p = 1, \dots, T \quad (4)$$

A fundamental question that arises is whether an equilibrium exists for the FT marketplace. The following theorem shows that such an equilibrium always exists.

THEOREM 2. (Existence of Nash-Equilibrium for FT Bandwidth Allocation) *There exists a set of per-user allocation vectors that, when feasible for each user, maximizes the total per-user allocation and is a NE.*

In order to prove this theorem, we prove the following lemma first.

LEMMA 1. (Existence and uniqueness of the per-user solution) *When the per-user FT maximization problem is feasible, there is a unique globally optimal solution (for a given set of allocations by the other agents).*

<sup>6</sup>Indeed, non-linear cost functions (of which ours is an instance) have been used before [12] to control congestion and achieve "proportional fairness."

PROOF. (*Sketch*) If the cost  $c_i < B_i$  when  $w_{ij} = 0$ , then there are feasible allocations of the fluid components  $w_{ij}$ . Notice also, that the feasible space defined as

$$\{w = (w_{ij}) \in \mathbb{R}^T | w_{ij} \geq 0 \text{ for } j = 1, \dots, T \text{ and } c_i \leq B_i\}$$

is convex. This follows from the fact that the constraints of equations 4 and 3 are concave functions. Then, by the Khun and Tucker (KT) theorem under convexity<sup>7</sup> there is vector  $w^*$  that maximizes the objective function. The uniqueness of this solution can be easily proven using the fact that the constraint 3 is strictly concave and the feasible space is convex.  $\square$

PROOF. *of Theorem 2* Define the following global fluid maximization problem:

Maximize

$$\sum_{i=1}^n \sum_{p=1}^T w_{ip} \quad (5)$$

subject to

$$c_i \leq B_i \quad \text{for } i = 1, \dots, n \quad (6)$$

$$w_{ip} \geq 0 \quad \text{for } i = 1, \dots, n \text{ and } p = 1, \dots, T \quad (7)$$

The Lagrangean of this problem is

$$L(w, \lambda, \gamma) = \sum_{i=1}^n \left( \sum_{p=1}^T w_{ip} + \sum_{p=1}^T \lambda_{ip} w_{ip} + \gamma_i c_i \right), \quad (8)$$

where  $w$  is the concatenation of the per-user allocation vectors, and  $\lambda_{ip}, \gamma_i$  are the Lagrange multipliers. Observe that eq. 8 is the sum of the corresponding Lagrangeans for the user problems, therefore a feasible  $w^*$  that maximizes 5 is also a global maximum for the per-user problems. Since the per-user allocations define a global maximum, no agent can improve its own objective by unilaterally deviating from this allocation vector, hence  $w^*$  is a NE.  $\square$

## 5. IMPLEMENTATION OF A T&C DSLAM MARKETPLACE

We describe a distributed implementation of the T&C marketplace, where there is one provider agent (running at the DSLAM for example), and a client-side agent running on the customer’s local router. The general architecture of the system is illustrated in Figure 2. In this architecture, the client-side agent is responsible for: (1) profiling the customer’s RT demand, (2) bidding for allocations during the bandwidth trading phase, and (3) shaping applications’ traffic according to reserved allocations. The provider-side agent provides two functionalities: (1) it runs the marketplace phases – bandwidth trading and bandwidth capping – just before the start of each epoch; and (2) once the epoch starts, enforces the allocations settled upon by the marketplace agents by using a traffic shaper for each customer line. The traffic shaper on the provider side enforces the total allocation determined by the T&C marketplace, but does not need to classify or monitor traffic, thus the overhead is minimal and the provider adheres to the principle of “net neutrality.”

The traffic shapers – both on the client-side and the provider-side – need not to be strict reservation-based shapers. The

<sup>7</sup>See theorem 7.16 [30]




Figure 2: Overall T&C Architecture

drawback of a strict reservation system is that it does not take advantage of the statistical multiplexing between the flows sharing the link. To avoid this limitation, we use a work-conserving scheduler, namely a derivative of the hierarchical link-sharing scheduler [11] – the Hierarchical Token Bucket (HTB) – which is currently available in the Linux kernel [7]. When using a work conserving scheduler, if some of the sources are idle, the unused capacity is distributed between the other sources. As a consequence, the reservations established in the T&C marketplace are minimum guarantees, but the aggregate utilization can always reach the total reserved capacity.

## 6. EXPERIMENTAL EVALUATION

In this section we use trace-driven simulations to (1) highlight the benefits that a user in our system begets by exhibiting some flexibility in scheduling its RT requests under T&C, (2) demonstrate the gains that an ISP stands to realize as a result of the overall smoother traffic profile of T&C, and (3) illustrate how various parameters affect the performance of T&C.

**Traces and Trace Pre-Processing:** As an alternative to direct DSLAM traces (which unfortunately were not available), we used publicly available WAN traces [25] to extract a slice of traffic associated with a customer access network. Table 1 shows the main characteristics of these WAN traces. During the pre-processing phase, the traffic from each individual IP address (assumed to be an end-user) is classified as RT or FT. RT in turn is decomposed into sessions, defined by sequences of contiguous non-zero RT demands. T&C operates by letting user agents express their flexibility or willingness to move RT sessions (forward or backward in time) some number of time slots. We define the *slack* as the number of time slots that an agent is allowed to shift an RT

Period	2009-03-31 00:00 to 2009-03-31 23:59
Total packets	1,551,089,845
TCP packets	1,194,409,653
UDP packets	4,321,852
Total TCP bytes (payload)	924,540,189,060

Table 1: Characteristics of the WAN trace used in our evaluation.

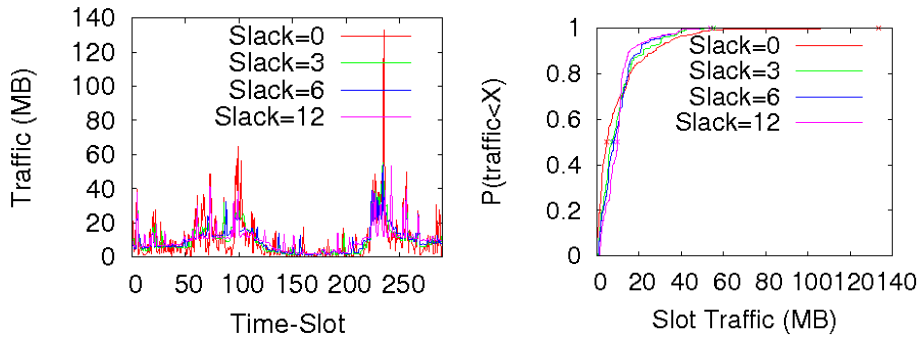


Figure 3: Utilization over time for RT requests with various slack values.

session (back or forth in time). A slack of 0 implies no flexibility. A slack of 1 implies an ability to shift an RT session by 5 minutes (our time slot) back or forth, if such a shift is advantageous. While the software agent acting on behalf of a user may use default slack settings (*e.g.*, depending on time of day, or type of application), we view this slack as a setting that users are able to fine-tune and/or adjust over time.

**How Does T&C Impact the ISP’s Bottom Line?** Our first experiment aims to evaluate how the 95th percentile of the ISP’s 5-minute traffic volume (the 95% traffic envelop) changes as a result of letting users schedule their RT requests according to the trading phase of T&C. For simplicity, we assume that all user agents adopt the same *slack* value for all RT sessions.

Figure 3 shows one example of the outcome after the market reaches an equilibrium.<sup>8</sup> On the left is the traffic per time-slot, and on the right is the CDF of traffic per time-slot. Table 2 shows the values of the 95% traffic envelop. These results underscore that selfishly scheduling RT requests yields an equilibrium with *significant* reduction in the 95% traffic envelop – up to 31% reduction when slack is 1 hours. Even for a small slack of 15 minutes, the savings amount to 16%.

We emphasize that the benefit from bandwidth trading quantified in the results in Table 2 (and elsewhere in this paper) is rather conservative given the nature of the WAN traces used in our evaluation, in which the peak-to-valley ratio is much lower than those observed in most characterization studies, *e.g.*, [20]. With workloads exhibiting typical variability, the benefits are likely to be even more significant.

Slack	95%	Savings%
0	36.3	0.0
3	30.6	15.6
6	27.4	24.4
12	24.9	31.4

Table 2: 95% traffic utilization (in MB) resulting from bandwidth trading under different slack values.

We now consider experiments in which both phases of T&C

<sup>8</sup>As detailed in the expanded version of this paper [21], reaching such an equilibrium based on best-response dynamics is practical for settings with hundreds of agents.




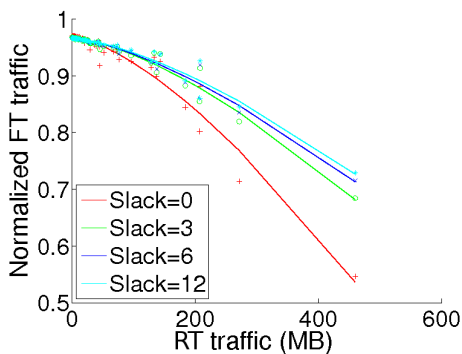
Figure 4: RT+FT traffic for various slack values.

are carried out. In particular, after completing the trading phase – thus scheduling all RT requests in the trace – user agents allocate as much fluid traffic as possible in accordance with their remaining budgets. Thus, an important consideration in setting-up these experiments is the budget assignment. In particular, we used the following policy: Let  $C$  denote the nominal traffic per time-slot that results in a total volume equal to the total traffic originally in the trace. We set the budget per customer to  $B_i = CT/n$ , which would yield the same overall traffic under an ideal allocation. Figure 4 shows the outcome of the two phases of T&C for various slack values. The y-axis is normalized with respect to  $C$ . Table 3 shows the 95% and 50% (median) of the time-slot utilizations, as well as the ratio between them. These results suggest that with T&C in place, the ratio is nearly 1.0, resulting in a perfect flattening of traffic over time slots, thus eliminating cost problems derived from spikes when using the 95/5 rule. Figure 4 also shows that the overall FT allocated to the users increases for increasing values of slack. This supports the idea that the marketplace provides the incentive for users to reveal the flexibility of their traffic demands.

**How Does T&C Impact the User’s Bottom Line?** To

	95%	Median	Ratio
Original	197.15	124.56	1.583
T&C	136.52	135.93	1.004

Table 3: Impact of T&C on traffic volume (in MB).



**Figure 5: RT and FT allocations per user for different slack values.**

evaluate T&C on a per-user basis, we compare how RT and FT allocations vary across users. Figure 5 shows a clear negative correlation between the allotment of FT and RT bandwidth. The relationship is not monotonic or deterministic because it depends on the outcomes from the trading phase, which affects the left-over budget for each agent. It is always the case though that the larger the slack afforded by users, the larger the FT allocation (points along the same vertical line in the plot).

## 7. RELATED WORK

While the application of game-theoretic and micro-economic approaches to networking problems is not novel [10, 12, 17, 22, 24], our approach of strategically trading-off allocation slots based on desirable properties for different traffic classes is new and quite promising.

Laoutaris and Rodriguez [19] recognized that the problems associated with rampant delay-tolerant traffic are due to the lack of incentives for end-users to properly schedule their delay-tolerant traffic and the lack of network mechanisms to identify and handle such traffic. As a solution to the first problem, they suggest giving users “higher-than-purchased” access rate during off-peak hours as a reward for time-shifting their delay-tolerant traffic. As a solution to the second problem, they propose the introduction of a store-and-forward service to handle the network transfer of bulk FT data during off-peak hours. The T&C marketplace is a mechanism that realizes the idea of providing incentives for the users to time-shift the traffic associated with their delay-tolerant applications. In a different setting, when a single authority has the control of the various applications sharing the resources, and can enforce an optimally computed schedule was presented by Laoutaris *et al* [20].

Recently, Briscoe *et al* [4] proposed an architecture that operates at the network edges and realizes the *cost fairness* model without directly charging users (hence, compatible with flat pricing). This work introduces *re-feedback*, a mechanism that allows measurement of downstream path metrics, such as delay and congestion. This information can then be used to police the compliance of end-users with a predetermined policy (*e.g.* backoff the sending rate in case of congestion). The network itself can perform the policing function requiring only a shaper at the ingress point and a dropper at the egress point. When doing so, it is the dominant strat-

egy for end-points to report the correct metrics. This is a congestion control mechanism, that provides the necessary feedback for flows to adjust their rates, and for the network to police response to congestion. It is strictly a best-effort scheme, and unlike T&C it does not provide the means for applications with specific Quality of Service (QoS) goals to make trade-offs that satisfy their requirements.

Approaches for *congestion-pricing* with explicit payments have been considered in a number of studies. Henderson *et al* [14] present a review of the benefits and limitations of these proposals. Examples include *Smart Markets* [22] and *Split-Edge Pricing* [3]. Of particular interest is the scheme proposed by Ganesh *et al* [12], which assigns costs to packets depending on congestion. Under a family of non-linear cost functions that depend on the utilization of the congested link and the flow’s demand, they showed convergence to steady-state equilibrium. While our mechanism and system model are entirely different, our cost function has similar characteristics.

Marbach [24] analyzes a priority queueing scheme (*a la* Diff-serv) where packets get charged based on their priority, and selfish users compete for bandwidth. Among other things, he shows that such a scheme leads to a Wardrop equilibrium and that allocation does not depend on the prices of each traffic class. A fundamental distinction in this case is that T&C enables different valuations for different classes of traffic, and uses these valuations to leverage the trading system.

A fundamental distinction between T&C and the various congestion pricing schemes considered in the literature ([4, 18, 14, 12]) is that none of these schemes takes into account the dual nature of RT versus FT applications. Therefore, all these schemes impose penalties (*e.g.* larger cost, increased drop rates) to all the traffic from a user during congestion periods. Because they operate over short-time-scales (targeting an instantaneous response to congestion), none of these approaches exploits the extra degree of freedom offered by the possibility of time-shifting RT allocations, or controlling the bandwidth consumption of FT applications.

## 8. CONCLUSION

Trade & Cap is an effective bandwidth management mechanism that enables software agents acting on behalf of self-interested users to collectively converge on an equitable allocation, based on the individual, private user valuation of network utility (*e.g.*, raw volume versus QoS over time).

T&C not only benefits users by allowing them to extract better utility from the network, but also it benefits the ISP by yielding smoother aggregate traffic volumes, which lowers traffic transit costs and reduces the currently unsustainable pressure on ISPs to upgrade their networks in order to keep up with peak demand. Under T&C, rather than acting as an arbiter, an ISP acts as an enforcer of what the community of rational users (sharing the contended resource) decides is a fair allocation of that resource. This is a welcome departure from current practices that force ISPs to use artificial notions of fairness to police shared bandwidth use, with negative implications to privacy and network neutrality.

## Acknowledgements

Jorge Londoño is supported in part by the Universidad Pontificia Bolivariana and COLCIENCIAS–Instituto Colombiano para el Desarrollo de la Ciencia y la Tecnología “Francisco José de Caldas”. Azer Bestavros is supported in part by NSF awards CNS-1012798, CNS-0952145, CCF-0820138, CSR-0720604, and EFRI-0735974.

## 9. REFERENCES

- [1] BERNAILLE, L., TEIXEIRA, R., AND SALAMATIAN, K. Early application identification. In *CoNEXT '06: Proceedings of the 2006 ACM CoNEXT conference* (New York, NY, USA, 2006), ACM, pp. 1–12.
- [2] BOSWORTH, M. H. Time warner: Metered broadband will prevent “internet brownouts”, April 10 2009.
- [3] BRISCOE, B. The direction of value flow in connectionless networks. In *Networked Group Communication* (1999), pp. 244–269.
- [4] BRISCOE, B., JACQUET, A., CAIRANO-GILFEDDER, C. D., SALVATORI, A., SOPPERA, A., AND KOYABE, M. Policing congestion response in an internetwork using re-feedback. *ACM SIGCOMM Computer Communication Review* 35, 4 (Aug 2005), 277–288.
- [5] CLARK, D. D., WROCLAWSKI, J., SOLLINS, K. R., AND BRADEN, R. Tussle in cyberspace: defining tomorrow’s internet. In *SIGCOMM* (New York, NY, USA, 2002), ACM, pp. 347–356.
- [6] CROWCROFT, J. Net neutrality: the technical side of the debate: a white paper. *SIGCOMM Computer Communications Review* 37, 1 (Jan 2007), 49–56.
- [7] DEVERA, M. HTB Home. Web page, July 2003.
- [8] DISCHINGER, M., HAEBERLEN, A., GUMMADI, K. P., AND SAROIU, S. Characterizing residential broadband networks. In *IMC '07: Proceedings of the 7th ACM SIGCOMM conference on Internet measurement* (New York, NY, USA, 2007), ACM, pp. 43–56.
- [9] F5 NETWORKS, INC. Bandwidth management for P2P applications.
- [10] FEIGENBAUM, J., PAPADIMITRIOU, C., SAMI, R., AND SHENKER, S. A BGP-based mechanism for lowest-cost routing. *Distrib. Comput.* 18, 1 (2005), 61–72.
- [11] FLOYD, S., AND JACOBSON, V. Link-sharing and resource management models for packet networks. *IEEE/ACM Trans. Netw.* 3, 4 (Aug. 1995), 365–386.
- [12] GANESH, A., LAEVENS, K., AND STEINBERG, R. Congestion pricing and user adaptation. In *IEEE INFOCOM* (2001), pp. 959–965.
- [13] GRUENER, W. Time warner shelves metered internet plans - for now, April 16 2009.
- [14] HENDERSON, T., CROWCROFT, J., AND BHATTI, S. Congestion pricing. Paying your way in communication networks. *IEEE Internet Comput.* 5, 5 (Sep/Oct 2001), 85–89.
- [15] IPOQUE. Bandwidth management with deep packet inspection, 2009.
- [16] KARAGIANNIS, T., PAPAGIANNAKI, D., AND FALOUTSOS, M. BLINC: Multilevel traffic classification in the dark. In *SIGCOMM* (2005).
- [17] KELLY, F. Charging and rate control for elastic traffic. *European Transactions on Telecommunications* 8 (1997), 33–37.
- [18] KELLY, F. P., MAULLOO, A., AND TAN, D. Rate control in communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research Society* 49 (1998), 237–252.
- [19] LAOUTARIS, N., AND RODRIGUEZ, P. Good Things Come to Those Who (Can) Wait or how to handle Delay Tolerant traffic and make peace on the Internet. In *HotNets'08* (2008).
- [20] LAOUTARIS, N., SMARAGDAKIS, G., RODRIGUEZ, P., AND SUNDARAM, R. Delay tolerant bulk data transfers on the internet. In *SIGMETRICS* (New York, NY, USA, 2009), ACM, pp. 229–238.
- [21] LONDOÑO, J., BESTAVROS, A., AND LAOUTARIS, N. Trade & cap: A customer-managed, market-based system for trading bandwidth allowances at a shared link. Tech. Rep. BUCS-TR-2009-25, Boston University, July 2009.
- [22] MACKIE-MASON, J., AND VARIAN, H. *Public Access to the Internet*. MIT Press, 1995, ch. Pricing the Internet.
- [23] MAIER, G., FELDMANN, A., PAXSON, V., AND ALLMAN, M. On dominant characteristics of residential broadband internet traffic. In *IMC '09: Proceedings of the 9th ACM SIGCOMM conference on Internet measurement conference* (New York, NY, USA, 2009), ACM, pp. 90–102.
- [24] MARBACH, P. Analysis of a static pricing scheme for priority services. *IEEE/ACM Trans. Netw.* 12, 2 (Apr 2004), 312–325.
- [25] MAWI WORKING GROUP. Traffic archive, 2009.
- [26] MOORE, A. W., AND ZUEV, D. Internet traffic classification using bayesian analysis techniques. In *SIGMETRICS* (New York, NY, USA, 2005), ACM, pp. 50–60.
- [27] NETO, H. T. M., ALMEIDA, J. M., ROCHA, L. C. D., MEIRA, W., GUERRA, P. H. C., AND ALMEIDA, V. A. F. A characterization of broadband user behavior and their e-business activities. *SIGMETRICS Perform. Eval. Rev.* 32, 3 (Dec. 2004), 3–13.
- [28] ROBERTS, L. G. A radical new router. *IEEE Spectr.* 46, 7 (July 2009), 34–39.
- [29] SEN, S., SPATSCHECK, O., AND WANG, D. Accurate, scalable in-network identification of P2P traffic using application signatures. In *WWW '04: Proceedings of the 13th international conference on World Wide Web* (New York, NY, USA, 2004), ACM, pp. 512–521.
- [30] SUNDARAM, R. K. *A First Course in Optimization Theory*. Cambridge University Press, 1996.