

An Integrated, Distributed Traffic Control Strategy for the Future Internet

Hao Che^a, Wenjing Su^b, Constantino Lagoa^b, Ke Xu^c, Chunyu Liu^c, and Yong Cui^c

^aThe Department of Computer Science and Engineering, the University of Texas at Arlington, TX, USA

^bThe Department of Electrical Engineering, Pennsylvania State University, PA, USA

^cThe Department of Computer Science, Tsinghua University, Beijing, China

Email: hche@cse.uta.edu, wus113@psu.edu, lagoa@engr.psu.edu, xuke@mail.tsinghua.edu.cn, liuchunyu@csnet1.cs.tsinghua.edu.cn, cy@csnet1.cs.tsinghua.edu.cn

Abstract¹

Due to the lack of a general theoretical foundation, today's distributed traffic control mechanisms developed at the networking layer, transport layer, and overlay are largely disintegrated. As a result, traffic control protocols developed at different layers may achieve conflicting design objectives and interact with one another in an unpredictable fashion. In this paper, we propose a novel strategy to tackle this issue. First, we propose a theoretical foundation for distributed traffic control. On the basis of this foundation, we then propose an integrated, multi-layer, multi-domain traffic control structure. This structure makes it possible to develop traffic control protocols at different layers, possessing the following nice features: (1) they achieve non-conflicting design objectives; (2) they enable rich service quality features, including Quality-of-Service (QoS), Traffic Engineering (TE), and Fast Failure Recovery (FFR); (3) they lead to highly scalable, globally stable and optimal control; (4) they can deal with network diversities and tussles among administrative domains; (5) they allow effective control of dynamically generated overlay networks. The proposed strategy only makes two assumptions about the Internet architecture, i.e., the ability to support multiple domains and multi-path forwarding. As a result, the proposed strategy can be applied to the existing or any future Internet architectures for which these two assumptions hold.

Keywords: Distributed Traffic Control, Traffic Engineering, Fast Failure Recovery, Quality of Service

1. Introduction

A. Background: The Internet by design provides only a single Best-Effort (BE) service. However, as it has evolved into a global commercial infrastructure, there has been an increasing demand for the Internet to provide rich service quality features to support sophisticated applications at global scale, including QoS, TE, and FFR. Enabling these features involves the development of distributed traffic control protocols that can quickly react to network congestions and link/node failures. As a result, various solutions were developed at different layers independently, such as distributed, multi-path enabled TE and FFR at the networking layer (e.g., TeXCP [5] and Mate [2]) and overlay (e.g., RON [1] and OverQoS [14]), and protocols at the transport layer with various degrees of involvement of network nodes (e.g., Fast TCP [4], VCP [15], and XCP [6]). However, "patching" the Internet with add-on service quality features at various layers

independently using diversified design methodologies is problematic for the following reasons.

First, without a sense of coordination, *the existing solutions developed at different layers often attempt to achieve conflicting design objectives*. For example, the existing TE mechanisms developed at the intra-domain level may adversely interact with the transport layer congestion control mechanisms, such as TCP window-based congestion control, in an unpredictable fashion. As another example, due to the lack of information with respect to the underlying resource availability and topology, the existing overlay traffic control mechanisms based on resource probing and topology discovery [1] [14] can adversely interact with the underlying networking layer traffic control mechanisms, resulting in unpredictable services at both layers, as observed in [10][13]. Moreover, most existing distributed traffic control protocols are empirically designed without a global design objective and provable properties such as stability.

Second, *the design spaces of the existing optimization-based approaches are too small to serve as theoretical underpinnings for the design of robust distributed traffic control protocols*. First, the existing approaches are exclusively rate-adaptation based and cannot provide rich QoS features. Second, the existing approaches that can support multi-path load balancing for TE and FFR work for a very limited number of design objectives or fairness criteria, such as min-max (e.g., [5]) or proportional fairness (e.g., [3]), inadequate for the support of multiple Classes of Service (CoSs). Third, the existing approaches do not have the flexibility to allow different degrees of engagement of the core nodes for distributed control. More specifically, different approaches may require different degrees of assistance from the core nodes, such as purely end-to-end [4], explicit-notification-based [15], or explicit-rate-based [6] traffic control at the transport layer. There has been no existing approach that allows various degrees of engagement of the core nodes. This makes it difficult to deal with network diversities and explore scalability-performance tradeoffs.

Apparently, "patching" the Internet with add-on service quality features at various layers independently has little chance to make the Internet a sustainable infrastructure. We believe that a clean slate, robust, scalable, integrated traffic control strategy must be sought, with minimum assumptions made about the networking architecture and its capabilities.

B. Proposed Strategy: The proposed strategy addresses the above issues by developing a *clean-slate, multi-domain, multi-layer, distributed traffic control structure*. This structure is based on a well-grounded theoretical foundation. More specifically, it relies on an optimization-based, distributed, QoS-aware, multi-path forwarding paradigm. In this forwarding paradigm, a set of distributed controllers are designed and applied to communication endpoints (i.e., end-hosts or servers), Internet access points, and domain edge nodes to enable QoS, TE, and FFR features end-to-end. The set of controllers in endpoints enable QoS-aware, end-to-end traffic control without any assistance from the network nodes. The set of controllers in Internet access points enables CoS features for flow aggregates sent to other Internet access points across the Internet. Internet domain edge nodes perform "per-hop" load balancing and/or rate adaptation for CoS-agnostic,

¹ Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

destination-based packet flow aggregates, enabling TE and FFR features. Here a next-hop node is another edge node, meaning that if a “next-hop” node is an edge node in the same domain as the traffic control edge node, it may actually be multiple hops away across the domain. As a result, load balancing and rate adaptation are performed among multiple next-hop nodes at the inter-domain level and among multiple paths from an ingress node to an egress node at the intra-domain level. As we shall discuss in detail in Section 3, under this structure, it becomes possible to develop traffic control protocols at the networking, transport, and overlay layers, possessing the following features:

1. They are integrated, achieving non-conflict design objectives;
2. they provide rich service quality features, including QoS, TE, and FFR;
3. They enjoy provable properties such as scalability, stability, and optimality;
4. They can deal with network diversities and tussles;
5. They allow effective control of dynamically generated overlay networks.

Finally, we note that the proposed structure is developed on the basis of a theoretical framework which is built on a continuous-time fluid-flow model and the above expected features are derived from the provable properties of the structure. It is a work-in-progress and overlooks some packet and session level effects including packet queuing, finite session lifetime, and granularity of control in both time and rate. As a result, how those effects will affect the protocol design in achieving the above desired features is subject to future investigation. Moreover, this paper only focuses on developing a theoretical underpinning and structure for protocol development, rather than protocols themselves.

The rest of the paper is organized as follows. Section 2 presents a general theoretical foundation for distributed traffic control. Section 3 describes the integrated traffic control structure. Finally, Section 4 concludes the paper.

2. Theoretical Foundation

Recently, we found, using the Sliding Mode technique in control theory, large families of controllers [8] [9] [11] [12], known as Optimization-based Distributed Control Laws (ODCLs). These families of ODCLs have overcome most of the limitations suffered by the existing distributed controllers, allowing distributed, multi-path, multi-CoS load balancing and/or rate adaptation for TE and FFR and large design space exploitation. However, like the existing solutions, they are *single-domain* solutions. In this section, we briefly describe a new family of ODCLs that overcomes this limitation and can scale to a large number of domains (the interested reader is referred to [16] for detailed description and proof of this family of ODCLs). Combined with the family of ODCLs described in [11], this family of ODCLs provides the much needed theoretical underpinning for the development of integrated traffic control structure.

Consider a multi-domain internet with access points attached to various internet domain edge nodes, serving as traffic source and sink nodes. Define a flow rate of a given type x_i ($i=1, 2, \dots, n$) as a CoS-based flow aggregate between a source access point and a sink access point and $x_{i,j}$ the sub-flow rate of x_i taking path j ($j=1, 2, \dots, n_j$). The objective is to find the allocation of traffic that leads to the maximization of the sum of individual flow utility functions:

$$J(x) = \sum_{i=1}^n U_i(x_i) = \sum_{i=1}^n U_i(x_{i,1}, x_{i,2}, \dots, x_{i,n_i})$$

subject to network link bandwidth constraints and QoS requirements.

Similar to the families of ODCLs in [8] [9] [11] [12], the new family of ODCLs that solves the above problem only requires that $U_i(x_{i,1}, x_{i,2}, \dots, x_{i,n_i})$ be a concave function with respect to each of its variables. It works as follows. A set of ODCLs in this family runs at both access points and domain edge nodes. They perform two level controls. At the inter-domain level, they perform load balancing/rate adaptation among multiple next-hop nodes. At the intra-domain level, they perform load balancing/rate adaptation among multiple paths leading to the same egress edge node. In what follows, such an edge-to-edge multi-path is called a *virtual link*. Hence, logically, one can view this approach as a hop-by-hop distributed traffic control scheme among edge nodes interconnected by physical and virtual links. The ODCL that controls the i -th flow running at the access points and edge nodes can be compactly written as follows:

$$\dot{x}_{i,j} = z_{i,j}(t, x_i, cg_i) \left[f_{i,j}(x_i) - (1 - \overline{cg}_{i,j} r_i^m r_i^M r_i^m r_i^{out}) \right] \quad (1)$$

$$\text{for } j = 1, 2, \dots, n_i,$$

where

$$f_{i,j}(x_i) = \begin{cases} 1 - e^{-\partial J / \partial x_{i,j}}, & \text{at an access point} \\ -1, & \text{at an edge node} \end{cases} ;$$

$$cg_{i,j} = \begin{cases} 1 & \text{if (virtual) link taken} \\ & \text{by } x_{i,j} \text{ is congested} \\ 0 & \text{otherwise} \end{cases}$$

and $\overline{cg}_{i,j}$ is the logical negation of $cg_{i,j}$; $z_{i,j}$ is a positive function that allows us to control the “speed” of the adaptation, which must satisfy a loose condition [16]; r_i^m is a parameter that may take two different values, determined by whether the sum of the received rates of the i -th flow from the upstream nodes is greater than the sum of the sending rates of the same flow to the downstream nodes ($r_i^m = 1$ for ODCLs running at a source access point); and r_i^{out} is a parameter that may take two different values, determined by whether the sum of the sending rates of the i -th flow to the downstream nodes is greater than the sum of receiving rates of the same flow at the downstream nodes ($r_i^{out} = 1$ for ODCLs running at the destination access point).

Now, the quantities r_i^m and r_i^M will allow one to address the case where different CoSs are to be enabled, including CoSs with a target rate, upper bounded rate, lower bounded rate, and both upper and lower bounded rates. *A CoS is enabled at the source access points only.* In other words, all the edge nodes only run ODCLs with $r_i^m = r_i^M = 1$ for all the flows traversing the

nodes. A right choice of these two parameters at the source access points enables a CoS-based ODCL for a given flow type (see [16] for more details). For example, if flows of type i are of the assured forwarding (AF) CoS, i.e., if they should achieve a given target rate $x_i = \Lambda_i$, then one should let $r_i^M = 1$ and,

$$r_i^m = \begin{cases} r_{\min}^{AF} < 1 & \text{if } x_i > \Lambda_i \\ r_{\max}^{AF} > 1 & \text{if } x_i < \Lambda_i \end{cases}$$

where r_{\min}^{AF} and r_{\max}^{AF} are constants. As another example, if flows of type i are of the BE CoS, then just let $r_i^m = r_i^M = 1$.

As proven in [16], the above family of ODCLs indeed leads to globally optimal and stable control in a multi-domain environment. Even in the presence of finite feedback delays, the flow rates will stay in the vicinity of the optimal ones and approach the optimal ones as the feedback delays diminish.

Note that the above family of ODCLs is fully distributed, in the sense that individual flows x_i 's are controlled by separate ODCLs, independently, at both access points and intermediate edge nodes the flows traverse. The interaction among flows is accounted for through a link or virtual link congestion indicator (i.e, $c\mathcal{G}_{i,j}$) only. However, one might note that the above

formulation requires that each edge node keep track of and run a separate ODCL for each and every flow x_i that traverses the node. This is not scalable because in this case, each domain edge node is required to keep the state information for each and every flow traversing it. To address this issue, we were able to show (see [16] for details) that the above family of ODCLs can be simplified to allow control of flows destined to the same destination network as a whole using one ODCL and no per flow state information is needed for the control. Moreover, this ODCL controls the *percentage* of rates sent to different paths/next-hops and there is no need to measure and control the individual flow rates. This reduces the needed ODCLs to be equal to the number of active IP prefixes, which can be effectively implemented in edge nodes with minimum software upgrades (see [12] for detailed discussion on how to implement this kind of ODCLs).

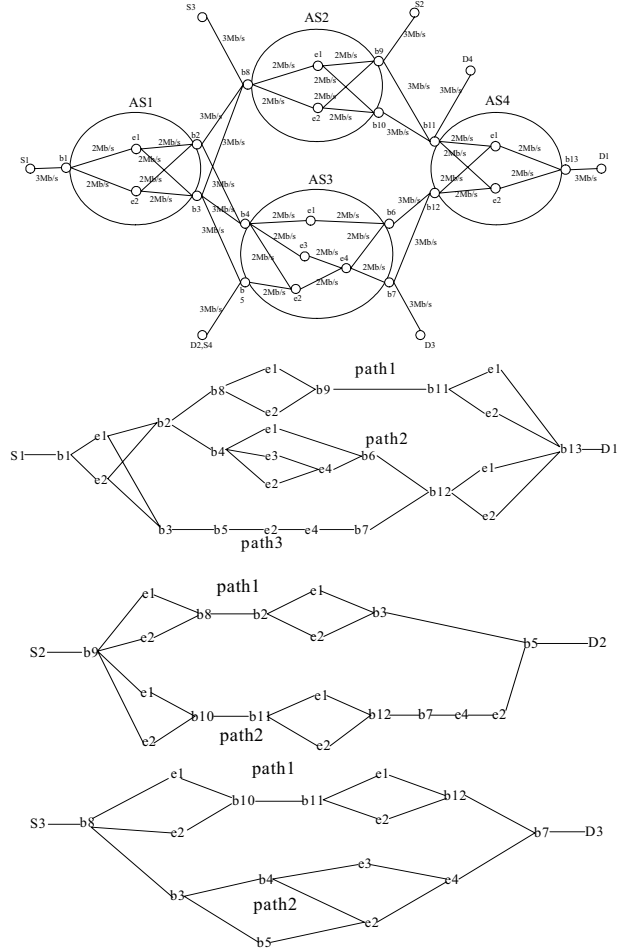
This family of ODCLs possesses the following three nice properties:

- (1) the CoS feature is enabled at the access points only. All the intermediate edge nodes perform CoS-agnostic, per-hop control to enable TE and FFR features.
- (2) it effectively breaks all the access-point-to-access-point control loops into a set of smaller, autonomous, per-hop control loops.
- (3) It works in a large design space and permits any degree of engagement of the core nodes for performance enhancement.

Property (1) is a consequence of use of CoS-agnostic ODCLs at all the edge nodes. Property (2) is due to the ability of the family of ODCLs to perform autonomous, coarse-granularity (i.e., per IP prefix based) control. The following two aspects of the family of ODCLs contribute to Property (3). First, since $U_i(x_i)$ is of a very general form, it allows a wide range of design objectives or fairness criteria to be exploited. Second, since $z_{i,j}$ can be any function in a large family of functions, a large number of ODCLs can be designed, corresponding to any given utility function or fairness criterion. In particular, by properly engineering $z_{i,j}$ to be a function of purely local information or information fed back from the core, the family of ODCLs permits any degree of involvement

of the core nodes for performance enhancement. This is made possible because at a minimum (e.g., letting $z_{i,j}$ be a function of purely local information or a constant), the family of ODCLs will work properly without the assistance from the core nodes. Note that in this case, the only non-local information needed in Eq (1) is $c\mathcal{G}_{i,j}$, which, however, can be inferred/probed by edge nodes without any assistance from the core nodes.

We also note that the family of ODCLs given in Eq. (1) degenerates to the family of ODCLs found in [11] by simply letting $r_i^{in} = r_i^{out} = 1$. This degenerated family of ODCLs runs between end-hosts only, taking the entire internet as a single domain. Since it retains all the nice properties of the above family of ODCLs, the degenerated family of ODCLs can serve as an ideal theoretical underpinning for the development of end-to-end protocols at the transport layer and host-based overlay. As shall be discussed in Section 3, combining these two families of ODCLs will allow us to develop a structure, under which integrated traffic control protocols at the networking, transport, and overlay layers can be developed.






Fig. 1 Network topology and multi-path for each source-destination pair

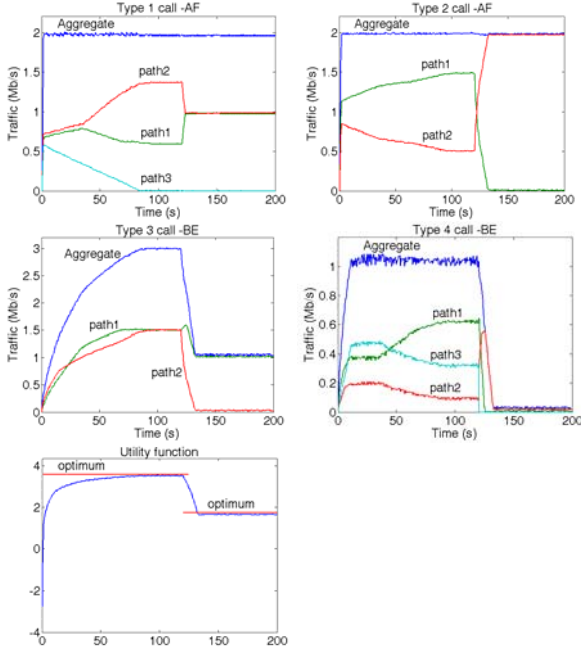


Fig. 2 Simulation results

Finally, the performance of the new family of ODCLs in Eq. (1) is tested based on a multi-domain scenario depicted in Fig. 1 (see [16] for more detailed simulation results). All the link bandwidths are indicated in the figure. Assume that there are four types of calls (or flows) corresponding to the source-destination pairs S_n - D_n , for $n = 1, 2, 3$, and 4. $n=1$ and 2 types of calls are AF CoS with target rate $\Lambda_1 = \Lambda_2 = 2\text{Mbps}$. $n=3$ and 4 are BE calls. Fig. 1 also depicts the multi-paths taken by these calls. All the calls take the same utility function of log form and z-function is chosen to be the same as the one selected in [12]. Figure 2 gives the simulation results for each type of call. At $t=120$ seconds, link $b3$ - $b5$ fails. Note that the data rate for each AF call converges very quickly to its target rate before and after the link failure, although sub-flows take longer time to stabilize (they “slide” over the optimal surface). The utility function converges to optimal values quickly even when the failure occurs.

3. Traffic Control Structure

On the basis of the two families of ODCLs presented in Section 2, now we develop an integrated traffic control structure in a multi-layer, multi-domain environment. The approach taken is to develop (1) an *end-to-end traffic control structure* that integrates host-based overlay traffic control with transport layer traffic control; (2) an *edge-to-edge traffic control structure* that integrates the networking layer traffic control with the network-based overlay traffic control; and (3) a *structure for end-to-end and edge-to-edge integration*. We describe these three structures, separately.

End-to-End Traffic Control Structure: Both native and host-based overlay applications have the same scope in terms of traffic control, i.e., end-to-end. This makes it possible to develop a common set of protocols based on the degenerated family of ODCLs, serving as the underlying transport for both native and host-based overlay applications. In terms of providing service quality features, our position is that end-to-end protocols should not be concerned with the TE and FFR features, but only the QoS feature. The TE and FFR features should be concerned by the networking-layer protocols at flow aggregate level. Hence, end-to-end traffic control protocols do not have to involve multi-path forwarding (as a result, we omit index j in the following discussion). Moreover, to be scalable, they should be purely end-to-end without engaging the network nodes for the control (i.e., z_{ij} should be a function of locally inferable information only).

With the degenerated family of ODCLs, a rich end-to-end QoS control structure can be exploited. The ODCLs with target rate guarantee may be used for the design of protocols to support applications with high throughput requirements. Note that the ODCLs with target rate guarantee can also pull back from their target rates gracefully if the network cannot sustain those rates. This nice feature makes it possible to design resilient differentiated services without complicated resource provisioning at the Internet edge. For example, one may consider the following utility function:

$$U(x) = \sum_{i=1}^n e_i^{AF} x_i^{AF} + \sum_{i=1}^n e_i^{BE} \log(x_i^{BE} + \beta)$$

where β is a constant, and e_i^{AF} and e_i^{BE} are the revenues earned by carrying one unit of flow i of AF CoS ($x_i^{AF} = \Lambda_i$) and one unit of $\log(x_i^{BE} + \beta)$ of the BE CoS, respectively. One can immediately derive two end-to-end ODCLs for the two CoSs, of which the BE CoS achieves the so called weighted proportional fairness [7]. By letting $z_i = \xi / (1 - e^{-\partial_j / \hat{\alpha}_i})$ for the BE CoS and applying the degenerated BE ODCL in Section 2, it can be easily shown that the resulting ODCL exhibits TCP-like behavior (see [11] for details). This will allow the interaction between the AF CoS and TCP to be studied. For example, using the earning ratio e_i^{AF} / e_i^{BE} as a tuning knob, the relative service quality for the two CoSs in terms of smoothness, responsiveness, and aggressiveness of control can be fine tuned to make the AF CoS a TCP-friendly service. Apparently, this degenerated family of ODCLs allows a broad design space to be exploited.

Edge-to-Edge Traffic Control Structure: A salient property of the families of ODCLs described in Section 2 is that they permit

various degrees of engagement of the core nodes. This makes it possible to design traffic control protocols at the network layer that permit different configurations corresponding to different degrees of involvement of the core nodes. This not only makes it possible to deal with network diversities, such as wired or wireless, and connectionless or connection-oriented, but also allows integration of traffic control at the network layer and network-based overlay. More specifically, the network-based overlay can simply adopt the same set of protocols used at the networking layer with configurations that do not engage core nodes. This ensures that overlay traffic control requires no knowledge about the network resource availability and topology², while allowing both layers to jointly achieve globally stable and optimal control. However, by doing so, care must be taken in defining the scope of domains. This is because a network-based overlay network may not coincide with a domain at the networking layer. In fact, an overlay network can span multiple networking-layer domains and may even partially overlap with a networking-layer domain. This issue can be addressed with the following integrated approach.

For traffic control purpose only, one may think of all the edge nodes of a networking-layer domain as overlay nodes and all the overlay nodes in the interior of a networking-layer domain as edge nodes of the network-layer domain in sinking and sourcing overlay traffic only. This effectively divides an overlay network across multiple networking-layer domains into non-overlapping overlay sub-domains, each coinciding with a networking-layer domain (with the edge of the domain augmented by the overlay nodes in the domain) and possibly other overlay sub-domains. This view will allow traffic control protocols based on the new family of ODCLs to be applied to the overlay networks as well as networking-layer networks in an integrated fashion. More specifically, service quality features can be enabled for any overlay network by simply turning on the same set of protocols running at the network layer with configurations that do not engage core nodes. This solution will eliminate any possible adverse interaction among overlay networks themselves, and between overlay and networking layer. Apparently, enabling/removing service quality features in an overlay is nothing more than activating/deactivating a set of protocols in a set of overlay nodes as well as a set of domain edge nodes at the networking layer.

Structure for End-to-end and Edge-to-edge Integration: The end-to-end and edge-to-edge traffic control can be integrated by mapping end-to-end QoS-based flows into a minimum of three CoSs enabled at each access point, including the BE, Assured Forwarding (AF) (with a target rate), and upper-bounded rate (UR) services.

The edge-to-edge BE and AF CoSs use the similar ODCLs as the end-to-end BE and AF CoSs (i.e., using the same utility function formats for conflict-free control). The only difference is that here the ODCLs control flow aggregates rather than individual application flows. The BE CoS will only be used for UDP-like non-rate-adaptive application flow aggregates without rate guarantee. The AF CoS will be used to support AF-based application flows with static call admission control to ensure that the sum of the target rates for application flows will not exceed the target rate of the AF CoS. The UR CoS will be used exclusively to support TCP-like adaptive flows without rate guarantee. The

² In the sense that it will adapt to whatever resources that are available and use whatever the routes found by a networking layer routing protocol.

UR CoS sets a bandwidth upper bound on the flow aggregate as a whole without limiting the number of adaptive flows that can be admitted. Unlike the existing networking layer adaptive traffic control schemes that can adversely interact with the end-to-end TCP congestion control, the UR CoS serves as a shield that protects the adaptive flows from being “adapted” again in the internet, but yet provide the needed TE and FFR features to, these flows.

Under the above structure, it becomes possible to develop integrated, service-quality-feature-rich, traffic control protocols that possess the following properties. First, *they enjoy provable properties including scalability, stability, and optimality*. Since the control is per-hop based involving endpoints, access points, and edge nodes only³, it is highly scalable and its implementation requires minimum software upgrade in these nodes. Moreover, it is proven that the control will be globally stable and optimal as long as the feedback delay for each per-hop control is within a given parameter range. This can be guaranteed by maintaining small domain sizes in terms of traffic control. Second, *they can deal with network diversities and tussles*. The ability to permit any degrees of engagement of the core nodes and autonomous control of a domain makes it possible to deal with network diversities, such as wired or wireless, connection-oriented or connectionless networks, and with or without sophisticated queuing mechanisms. The autonomous control at the inter-domain level also makes it possible to address tussles by selecting next-hop domains based on administrative policies, such as service level agreements, between domains. Third, *they allow effective control of dynamically generated overlay networks*. This is made possible because a domain edge node in terms of traffic control can be easily added/removed by simply activating/deactivating a set of protocols in that node.

Finally, we note that this paper is work-in-progress and issues pertaining to specific protocol design are not addressed. Moreover, the above expected features are derived from the provable properties of the structure, on the basis of a control theoretical framework. This theoretical framework, however, provides an approximate view of the real-world traffic by modeling it using a continuous-time fluid-flow model. It overlooks some packet and session level effects, such as packet queuing, finite session lifetime, granularity of control in both time and rate, etc. As a result, the proposed structure should be taken as a theoretical underpinning to guide the protocol development. How faithfully or closely the protocols developed under this structure will achieve the desired features in practice is subject to future investigation.

4. Conclusions

This paper developed an integrated traffic control structure in a multi-layer, multi-domain environment. It achieves the following design objectives: (1) the end-to-end control enables rich QoS features for application flows, whereas the edge-to-edge control enables TE and FFR features for flow aggregates across multiple domains; (2) traffic controls at all layers achieve mutually non-conflicting design objectives or jointly achieve a common design objective, with provable stability and optimality properties; (3) edge-to-edge traffic control at the intra-domain level can work

³ In the case when domain core nodes are engaged in the control, the complexity would be higher. This, however, is a user choice, meaning that each domain can choose to involve or not involve the core nodes for traffic control purposes.

using locally inferable information only, while allowing for various degrees of involvement of network core nodes for performance enhancement. This traffic control structure sets a foundation underlying the development of scalable and robust distributed traffic control protocols in an integrated fashion. This work is work-in-progress and issues pertaining to specific protocol design will be addressed in the future.

AKNOWLEDGEMENT: This work is supported by NSF under grants ANI-0125653 and NeTS-NBD 0519999.

REFERENCES

- [1] D. G. Andersen, H. Balakrishnan, M. F. Kaashoek, and R. Morris, "Resilient Overlay Network," in *Proceedings of ACM SOSP*, pp. 131- 145, Oct. 2001.
- [2] A. Elwalid, C. Jin, S. Low, and I. Widjaja, "MATE: MPLS Adaptive Traffic Engineering," in *Proceedings of IEEE INFOCOM*, 2001.
- [3] H. Han, S. Shakkottai, C. V. Hollot, R. Srikant, and D. Towsley, "Overlay TCP for Multi-Path Routing and Congestion Control," in *ENS-INRIA ARC-TCP Workshop*, Nov. 2003.
- [4] C. Jin, D. Wei, and S. Low, "FAST TCP: Motivation, Architecture, Algorithms, and Performance," in *Proceedings of IEEE INFOCOM*, March 2004.
- [5] S. Kandula, D. Katabi, B. Davie, and A. Charny, "TeXCP: Responsive yet Stable Traffic Engineering," in *Proceedings of ACM SIGCOMM*, 2005.
- [6] D. Katabi, M. Handley, and C. Rohrs, "Congestion Control for High Bandwidth-Delay Product Networks," in *Proceedings of ACM SIGCOMM*, 2002.
- [7] 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, no. 3, pp. 237-252, Mar. 1998.
- [8] C. Lagoa and H. Che, "Decentralized Optimal Traffic Engineering for the Internet," *ACM SIGCOMM Computer Communication Review*, vol. 3, no. 5, Oct. 2000.
- [9] C. Lagoa, H. Che, and B. Movsichoff, "Adaptive Control Algorithms for Decentralized Optimal Traffic Engineering," *IEEE/ACM Transactions on Networking*, vol. 12, no. 3, June 2004.
- [10] Y. Liu, H. Zhang, W. Gong, D. Towsley, "On the Interaction Between Overlay Routing and Underlying Routing," in *Proceedings of IEEE INFOCOM*, March 2005.
- [11] B. Movsichoff, C. Lagoa, and H. Che, "Minimal Feedback Optimal Algorithms for Traffic Engineering in Computer Networks," accepted for publication in *IEEE/ACM Transactions on Networking*, available at: <http://crystal.uta.edu/~hche/PUBLICATIONS/publication.htm>.
- [12] B. Movsichoff, C. Lagoa, and H. Che, "Decentralized Optimal Traffic Engineering in Connectionless Networks," *IEEE Journal on Selected Areas in Communications*, vol. 23, no. 2, pp. 293-303, Feb. 2005.
- [13] L. Qiu, Y. R. Yang, Y. Zhang, and S. Shenker, "On Selfish Routing in Internet-Like Environments," in *Proceedings of ACM SIGCOMM*, Aug. 2003.
- [14] L. Subramanian, I. Stoica, H. Balakrishnan, and R. H. Katz, "Over QoS: Offering Internet QoS Using Overlays," in *Proceedings of HotNet-I Workshop*, Oct. 2002.
- [15] Y. Xia, L. Subramanian, I. Stoica, and S. Kalyanaraman, "One More Bit is Enough," in *Proceedings of ACM SIGCOMM*, 2005.
- [16] W. Su, C. Liu, C. Lagoa, H. Che, K. Xu, and Y. Cui, "A Family of Optimization-based Distributed Traffic Control Laws in an Multi-domain Internet," Technical Report, available at: <http://crystal.uta.edu/~hche/PUBLICATIONS/publication.htm>.