Adaptive optimization for active queue management supporting TCP flows

Baldi, Simone; Kosmatopoulos, Elias B.; Pitsillides, Andreas; Lestas, Marios; Ioannou, Petros A.; Wan, Yiming

DOI

[10.1109/ACC.2016.7525004](https://doi.org/10.1109/ACC.2016.7525004)

Publication date

2016

Document Version

Accepted author manuscript

Published in

Proceedings of the 2016 American Control Conference (ACC2016)

Citation (APA)

Baldi, S., Kosmatopoulos, E. B., Pitsillides, A., Lestas, M., Ioannou, P. A., & Wan, Y. (2016). Adaptive optimization for active queue management supporting TCP flows. In G. Chiu, K. Johnson, & D. Abramovitch (Eds.), *Proceedings of the 2016 American Control Conference (ACC2016)* (pp. 751-756). IEEE .
<https://doi.org/10.1109/ACC.2016.7525004>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Adaptive Optimization for Active Queue Management supporting TCP Flows

Simone Baldi, Elias B. Kosmatopoulos, Andreas Pitsillides, Marios Lestas, Petros A. Ioannou, and Yiming Wan

Abstract—An adaptive decentralized strategy for active queue management of TCP flows over communication networks is presented. The proposed strategy solves locally, at each link, an optimal control problem, minimizing a cost composed of residual capacity and buffer queue size. The solution of the optimal control problem exploits an adaptive optimization algorithm aiming at adaptively minimizing a suitable approximation of the Hamilton-Jacobi-Bellman equation associated with the optimal control problem. Simulations results, obtained by using a fluid flow based model of the communication network and a common network topology, show improvement with respect to the Random Early Detection strategy. Besides, it is shown that the performance of the proposed decentralized solution is comparable with the performance obtained with a centralized strategy, which solves the optimal control problem via a central unit that maintains the flow states of the entire network.

I. INTRODUCTION

The Transmission Control Protocol (TCP), used in many communication networks and over the Internet, controls the rate at which packets are sent from sources across the network, using end-to-end implicit feedback, mainly through inferred packet loss. Due to the ever increasing use of the Internet, existing TCP congestion controls based on Jacobson [1] and its variants (e.g. Tahoe, Reno, New Reno, etc...) are increasingly becoming ineffective. Newer variants of the TCP protocol advocate the Active Queue Management (AQM) protocol responsible for marking/dropping packets at the router level (e.g. RED-Random Early Detection [2]) in order to signpost the state of congestion of the network early, or even provide explicit control feedback (Explicit Congestion Notification), based on the router state. The AQM schemes implemented in the routers of communication networks supporting TCP flows can be modeled as a feedback control system, and the problem of finding an efficient

marking strategy that minimizes network congestion is recast as a control design problem.

The design and control of modern communication networks raises several issues, the main one being the need for decentralized algorithms [3]: the challenge is to understand how such algorithms can be designed so that the network as a whole reacts intelligently to congestion.

Several mathematical models of AQM schemes supporting TCP flows in communication networks have been proposed in literature, e.g. [4], [5], [6]. These models are based on delay differential equations that can describe with different levels of complexity the behavior of the network system, according to the assumptions made (time-invariant/time-varying round trip time, etc). In contrast with discrete event simulators [7], fluid flow based models can be generalized to networks of arbitrary complexity and topology, and do not suffer from problems of scalability. Furthermore, fluid flow model offer a better insight for analysis and design of the control system. The stability analysis of fluid flow models for describing high-speed network behavior represents a subject of recurring interest in the last years [8], [9], [10], [11], [12].

The Random Early Detection [2] is one of the most popular AQM marking strategies: in RED algorithm, the price (packet marking/dropping mechanism) is probabilistically adjusted according to a nonlinear drop function depending on the average queue length at the router. The RED strategy relies on a few parameters and it is currently implemented in many communication networks thanks to its simplicity. However, it is well known that TCP/RED can oscillate wildly and it is extremely hard to reduce the oscillation by tuning RED parameters. Oscillatory behavior is mitigated but not avoided using heuristic modifications of the RED strategy, e.g. ARED [13] and FRED [14]. Many models have shown that oscillation is an inevitable outcome of the protocol itself [15], [4]. TCP/RED becomes unstable when delay increases, or more strikingly, when link capacity increases. The analysis illustrates the difficulty of setting RED parameters to stabilize TCP: they can be tuned to improve stability, but only at the cost of large queues [16]. In particular, it was shown that for increasing time delays and capacities, the TCP/RED fixed-point equilibrium becomes unstable, and the trajectories oscillate around a stable limit cycle. This suggests that the current RED protocol is not well-suited for future networks where capacity will be large. This is the reason why, over the years, many different protocols have been studied as an alternative to the RED strategy [17], [18], [19],

S. Baldi and Y. Wan are with the Delft Center of Systems and Control {s.baldi,y.wan}@tudelft.nl.

E. B. Kosmatopoulos is with Dept. of Electrical and Computer Engineering, Democritus University of Thrace, Xanthi, Greece and the Informatics & Telematics Institute, Center for Research and Technology Hellas (ITI-CERTH), Thessaloniki, Greece kosmatop@iti.gr.

A. Pitsillides is with Computer Science Department, University of Cyprus (UCY), Nicosia, Cyprus pitsillides@cs.ucy.ac.cy.

M. Lestas is with the Department of Electrical Engineering, Frederick University, Nicosia, Cyprus, eng.lm@frederick.ac.cy.

P. A. Ioannou is with Department of Electrical Engineering, University of Southern California, Los Angeles, CA, USA ioannou@usc.edu.

The research leading to these results has been partially funded by the European Commission FP7-ICT-5-3.5, Engineering of Networked Monitoring and Control Systems, under the contract #257806 (AGILE).

[20], [21], [22].

This paper proposes an adaptive optimization method for active queue management of TCP flows. The contribution of the proposed research lies in the paradigm of adaptive (approximately) optimal congestion control, where the objective is to approximately solve the Hamilton-Jacobi-Bellman equation associated to the optimal control problem. In this work, in order to take into account the decentralized structure of the network, the optimization problem is solved separately at each link, using only local information. The decentralized solution is compared with the (not implementable) centralized one, and it is demonstrated via simulations that the achieved performances are comparable.

The paper is organized as follows: Section II introduces the congestion control problem over communication networks supporting TCP flows. The optimal congestion control is formulated in Section III and the proposed adaptive optimization algorithm is presented in Section IV. Simulative results are given in Section V.

II. CONGESTION CONTROL FRAMEWORK

Fig. 1: Feedback representation of the congestion control problem

We consider an AQM-based communication network setting, with Explicit Congestion Notification (ECN). The network is modeled as a set of L links with finite capacities c_j , $j = 1, \dots, L$. The links are shared by a set of N sources indexed by $i = 1, \dots, N$. According to the network topology, not all sources use all the links. In general, each source i uses a subset L_i of the links, *i.e.*, $L_i \subseteq L$. The network topology is defined by the $L \times N$ routing matrix A . Each element A_{ji} of the matrix is defined as

$$A_{ji} = \begin{cases} 1 & \text{if } j \in L_i \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

Associated with each link j is its marking probability $p_j(t)$ at time t , and associated with each source i is its window $w_i(t)$ at time t (or, equivalently, its sending rate $s_i(t)$ at time t). The overall network can be described by the diagram of Fig. 1. Propagation delays must be taken into account. Define the round trip time of source i as $\tau_i(t)$: then the sending rate

of source i is

$$s_i(t) = \frac{w_i(t)}{\tau_i(t)}. \quad (2)$$

What is available at each link j is the aggregate flow rate, which is

$$y_j(t) = \sum_{i=1}^N A_{ji} s_i(t - \tau_{ji}^f(t)), \quad (3)$$

where τ_{ji}^f is the forward delay from source i to link j . There are two commonly adopted ways to define the end-to-end marking probability observed at source i . The first one is

$$q_i(t) = \sum_{j=1}^L A_{ji} p_j(t - \tau_{ji}^b(t)), \quad i = 1, \dots, N, \quad (4)$$

where τ_{ji}^b is the backward delay from link j to source i . Eq. (4) leads to an allocation of the sending rates characterized by the proportional fairness criterion. The second alternative, which will be the one considered in this paper is

$$q_i(t) = \min_{j \in L_i} \{A_{ji} p_j(t - \tau_{ji}^b(t))\}, \quad i = 1, \dots, N, \quad (5)$$

leading to an allocation of the sending rates characterized by the max-min fairness criterion. The forward and backward delays are related to the round trip time through

$$\tau_i(t) = \tau_{ji}^f(t) + \tau_{ji}^b(t). \quad (6)$$

Next, we model the source and the link behavior. The TCP Reno is chosen as the protocol describing the source behavior. The TCP model is obtained by ignoring the TCP slow start and time out mechanisms. It is well known that the congestion window size is increased by one every round trip time if no congestion is detected, and is halved upon a congestion detection. This additive-increase multiplicative-decrease behavior of TCP has been modeled in [4] by the following difference equation

$$\dot{w}_i(t) = \frac{s_i(t - \tau_i(t))(1 - q_i(t))}{w_i(t)} - \frac{s_i(t - \tau_i(t))q_i(t)w_i(t)}{2}. \quad (7)$$

The first term of (7) describes the window's additive increase phase, and the second term the multiplicative decreasing phase, according to the packet marking probability.

The link behavior is composed of the AQM dynamics, which describes the way the marking probability is adjusted, and of the queuing dynamics. The following queuing model is assumed

$$\dot{b}_j(t) = \begin{cases} y_j - c_j & \text{if } b_j > 0 \text{ or } y_j - c_j \geq 0 \\ 0 & \text{if } b_j \leq 0 \text{ and } y_j - c_j \leq 0 \end{cases}, \quad (8)$$

where c_j is the link capacity. For the moment the AQM dynamics are given by the RED strategy, which updates the marking probability according to the average queue length. The average queue length $r_j(t)$ is updated according to

$$\dot{r}_j(t) = -K_j c_j (r_j(t) - b_j(t)), \quad (9)$$

where $b_j(t)$ is the instantaneous queue length and $0 < K_j < 1$. The RED marking probability is given by

$$p_j(t) = \begin{cases} 0 & \text{if } 0 \leq r_j(t) < \text{min_th} \\ (r_j(t) - \text{min_th})\rho & \text{if } \text{min_th} \leq r_j(t) < \text{max_th} \\ 1 & \text{if } r_j(t) \geq \text{max_th} \end{cases} \quad (10)$$

where $\rho = p_{\max}/(\text{max_th} - \text{min_th})$, and $0 < p_{\max} < 1$. The parameters $(\text{min_th}, \text{max_th}, p_{\max}, K)$ are the configurable parameters of the RED strategy. A graphical representation of the RED drop function is given in Fig. 2.

Fig. 2: RED drop function

III. OPTIMAL CONGESTION CONTROL

Define $s = [s_1, s_2, \dots, s_N]'$ and $C = [c_1, c_2, \dots, c_L]'$. It is well known that the congestion problem can be seen as a maximization problem of a source utility function

$$\begin{aligned} \max_{s \geq 0} \quad & \sum_{i=1}^N U_i(s_i) \\ \text{s.t.} \quad & As \leq C, \end{aligned} \quad (11)$$

where for the described TCP Reno protocol the utility function $U_i(s_i)$ takes the form

$$U_i(s_i) = \frac{\sqrt{2}}{\tau_i} \tan^{-1} \left(\frac{\tau_i s_i}{\sqrt{2}} \right). \quad (12)$$

We recast the problem as a minimization problem by defining

$$\bar{U}_i(s_i) = \frac{\tau_i}{\sqrt{2}} \tan^{-1} \left(\frac{\sqrt{2}}{\tau_i s_i} \right). \quad (13)$$

Furthermore, assuming small delay, we use the Pade approximation

$$s(t - \tau_i) \approx \frac{1 + \tau_i/2}{1 - \tau_i/2} s(t) \quad (14)$$

to approximate the delay in (7). Finally, we introduce the low-pass filter

$$\dot{p}_i(t) = -\lambda p_i(t) + \lambda u_i(t), \quad (15)$$

with $\lambda > 0$. Taking the integral over time of (13), we end up with the minimization problem

$$\begin{aligned} \min \quad & \int_0^\infty \left[\sum_{i=1}^N \bar{U}_i(s_i(t)) \right] dt \\ \text{s.t.} \quad & As \leq C, s \geq 0 \\ & 0 \leq p \leq 1, \end{aligned} \quad (16)$$

which present two main problems. The first problem is the fact that the sending rate for each source is not measurable, so that the utility function for each source cannot be evaluated. To this purpose we replace (13) with

$$\min \int_0^\infty \sum_{j=1}^L [e_j^2(t) + b_j^2(t) + \rho (p_j^2(t) + u_j^2(t))] dt, \quad (17)$$

where $e_j = y_j - c_j$ is the residual capacity at link j , and $\rho > 0$ controls the control effort: (17) expresses the desire of matching the data flow to the link capacity while maintaining small queue sizes and small control efforts. The second problem is that minimizing (16) require maintenance of per flow states within the network, resulting in an infeasible centralized approach due to the distributed nature of the Internet. For this reason we consider at every link j the local cost

$$\min \int_0^\infty [e_j^2(t) + b_j^2(t) + \rho (p_j^2(t) + u_j^2(t))] dt \quad (18)$$

which we aim at minimizing locally at the link level.

The local dynamics, i.e. the dynamics at each link j , can be described by the following system:

$$\begin{aligned} \begin{bmatrix} \dot{s}(t) \\ \dot{d}_j(t) \\ \dot{b}_j(t) \\ \dot{q}_j(t) \end{bmatrix} &= \begin{bmatrix} f_{TCP}(d(t), p(t)) \\ f_{PADE}(s(t), d(t)) \\ f_{QUEUE}(s(t)) \\ -\lambda p(t) \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ 0 \\ \lambda \end{bmatrix} u_j(t) \\ \begin{bmatrix} y_j(t) \\ b_j(t) \\ p_j(t) \end{bmatrix} &= \begin{bmatrix} A(j,:) & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} s(t) \\ d_j(t) \\ b_j(t) \\ p_j(t) \end{bmatrix} \\ &= H \begin{bmatrix} s(t) \\ d_j(t) \\ b_j(t) \\ p_j(t) \end{bmatrix}, \end{aligned} \quad (19) \quad (20)$$

where $f_{TCP}(s(t), d(t), p(t))$, $f_{PADE}(s(t))$ and $f_{QUEUE}(s(t))$ and defined by the TCP-Reno dynamics (7), by the Pade approximation (14) and by the queuing dynamics (8), respectively. Note that $A(j,:)$ indicates the j -th row of the routing matrix A . Define the state $x(t) = [s'(t) \ d_j'(t) \ b_j'(t) \ p_j'(t)]'$: so, y_j might be considered as an output variable rather than a state variable. The consequence for this is that the optimal control formulation must be in an output-feedback form, rather than in a state-feedback one.

IV. ADAPTIVE OPTIMIZATION ALGORITHM

In this section, the RED strategy is substituted by a novel AQM strategy, based on the adaptive (approximate) solution of an optimal control problem (AO for brevity). By

introducing the constraints in (16) as penalty functions, we can rewrite the minimization problem (the local subscript j is omitted for simplicity) as:

$$J = \int_0^{\infty} [\Pi(x(t)) + u'(t)Ru(t)] dt, \quad (21)$$

where $R = \text{diag}(\rho)$, subject to the dynamics (19).

According to the Hamilton-Jacobi-Bellman (HJB) equation, the controller that optimizes the system performance (21) can be obtained as the solution of the following partial differential equation:

$$\dot{V}^*(x(t)) = \left(\frac{\partial V^*}{\partial x} \right) (f(x) + Bu^*) = -(\Pi(x) + u'Ru), \quad (22)$$

where $B = [0 \ 0 \ 0 \ \lambda]'$ and $f(x)$ come from (19). $V^*(x)$ denotes the so-called *optimal cost-to-go function* and u^* denotes the optimal controller, which can be seen to satisfy

$$u^* = -\frac{1}{2}R^{-1}B' \left(\frac{\partial V^*}{\partial x} \right)'. \quad (23)$$

The exact solution of (22) is computationally very expensive. For this reason, an approximated parametrized solution is taken to as follows

$$V(x) = z'(x)Pz(x), \quad u = -R^{-1}B'\Xi_z(x)Pz(x), \quad (24)$$

where $\Xi_z(x)$ is the Jacobian matrix of $z(x)$ with respect to x and

$$z(x) = \begin{bmatrix} \sqrt{\beta_1(x)}x \\ \sqrt{\beta_2(x)}x \end{bmatrix}, \quad (25)$$

$$P = \begin{bmatrix} P_1 & 0 \\ 0 & P_2 \end{bmatrix}, \quad (26)$$

where P_i are *positive definite matrices*. The functions $\beta_i(x)$ are the so-called *mixing signals* $\beta_i(x)$ [23] designed as

$$\beta_1(e_j, q) = \frac{1}{1 + e^{-he_j b_j}} \quad (27)$$

$$\beta_2(e_j, q) = 1 - \beta_1(e_j, b_j), \quad (28)$$

where $h > 0$ is a parameter which regulates the sharpness of the sigmoid. The functions in (27) are motivated by the switching behavior introduced by the queuing dynamics (8). The functions in (27) are smooth and give rise to a piecewise quadratic value function and a piecewise linear control law: the two functions are shown in Fig. 3.

By using the approximations (24) and integrating (22) in the interval $[t, t + \delta t]$, where $\delta t > 0$ is a discretization step, one can see that in case the optimal controller u^* were applied then,

$$\Delta V(x(t)) \approx - \int_t^{t+\delta t} [\Pi(x(r)) + u^{*'}(r)Ru^*(r)] dr, \quad (29)$$

where $\Delta V(x(t)) = V(x(t + \delta t)) - V(x(t))$. Having the above equation in mind and the provided approximations, let us assume that the following controller is applied to the actual system:

$$\hat{u} = \hat{u}(x(t); \hat{P}) = -R^{-1}B'\Xi_z(x)\hat{P}z(x), \quad (30)$$

Fig. 3: Activation functions: β_1 (lower), β_2 (upper)

where $\hat{P}$ denotes an estimate of the unknown matrix P . Let us also define the following “error” term

$$\varepsilon(x(t), \hat{P}) = \Delta \hat{V}(t) + \int_t^{t+\delta t} [\Pi(x(r)) + \hat{u}'(r)R\hat{u}(r)] dr, \quad (31)$$

where $\hat{V} = \hat{V}(x(t); \hat{P}) = z'(x)\hat{P}z(x)$ and $\Delta \hat{V}(t) = \hat{V}(x(t + \delta t)) - \hat{V}(x(t))$. By using equation (29) it can be seen that the “error” term $\varepsilon(x(t), \hat{P})$ provides us with a “measure” of how far the estimate $\hat{P}$ is from its optimal value P . Thus the following update law is proposed, based on the stochastic approximation algorithm [24]

$$\hat{P}_{t+\Delta t} = \text{Pr}[\hat{P}_t - \alpha(t)\nabla_{\hat{P}}\varepsilon^2(x(t), \hat{P}) + e(t) + \alpha(t)\xi(t)]_2$$

where the symbol $\text{Pr}[\cdot]$ denotes the operation of projecting $(\cdot)$ into the set of positive definite matrices, $e(t)$ is a term that exponentially decays to zero and $\xi(t)$ is a zero-mean random term. The term $\nabla_{\hat{P}}\varepsilon^2(x(t), \hat{P})$ is calculated from the knowledge of the system dynamics (19) (details are straightforward and omitted for lack of space). The update law (32) guarantees that $\hat{P}_t$ converges almost surely to the following set

$$\mathcal{E} = \{\hat{P} : \hat{P} \succ 0 \text{ and } \nabla_{\hat{P}}\varepsilon^2(x(t), \hat{P}) = 0\}. \quad (33)$$

Remark 1: The control law should be in an output-feedback form, rather than in a state-feedback one. The optimal static output-feedback control formulation proceeds along very similar steps as the optimal state-feedback control formulation, as can be seen for example in [25, Sect. 6.5], with the additional condition that the optimal cost-to-go function must satisfy a structural constraint. Such a constraint, using the proposed approximations, can be expressed in the following form:

$$V(x) = \sum_{i=1}^M \beta_i(x')P_i x = \sum_{i=1}^M \beta_i(x')H'P_i^c H x \quad (34)$$

with H as in (20). A problem associated with the structural constraint (34) is that it requires the knowledge of $A(j, \cdot)$, i.e., of the topology of the network. Such a knowledge, which can be easily assumed in a centralized architecture, is not available at a link level. For this reason the decentralized

solution has been obtained by assuming $A(j,:) = [1, \dots, 1]$, i.e., by assuming that every link is a bottleneck link. Such an approximation has been verified via simulations to lead to satisfactory results. Note that under such an approximation the problem can be even more simplified, since $y_j = Ns_i$, so that y_j can be considered as a state variable instead of s , and the optimal state-feedback formulation is completely recovered.

V. SIMULATION RESULTS

All the simulations are conducted using the fluid flow based model of Sect. II. Persistent FTP sources are considered, with the packet size equal to 1000 bytes. The decentralized solution, denoted for brevity decentralized AO, will be compared with the corresponding centralized one (implementable only in simulations), denoted for brevity centralized AO. The numerical results are obtained by using a typical network topology.

A single bottleneck link with capacity 9 pkts/ms is shared by 50 persistent FTP sources. The round trip propagation delay is 40 ms. The RED parameters has been tuned as in [16], i.e., $\max_p = 0.1$, $\min_th = 50$ pkts, $\max_th = 550$ pkts, and weight for queue averaging $K = 10^{-4}$.

Fig. 4: Single Bottleneck topology

Fig. 5 show the behavior of the centralized and of the decentralized AO (after convergence) as compared with the RED strategy. It can be seen that while the settling time of the three strategies is comparable, the AO controllers completely remove the oscillatory behavior of the RED strategy. Notice that all the three strategies start with a half-full buffer, but while the RED controller has oscillatory queuing behavior, the AO controllers empty the buffer and keeps its size at value which is almost zero. The difference between the RED and the AO strategies is also underlined by looking at the sending rate trajectory: while, in the transient phase, the RED controller increases the sending rate with respect to its initial value, the AO controllers decreases it. Table II summarizes the improvement in terms of the cost function (17): the performances of the centralized and of the decentralized AO solutions, with improvements over RED ranging from 67 to 80%. Moreover, also in the presence of larger capacity and round trip-time the proposed AO strategies sensibly reduce the oscillations and achieve smaller residual capacity and queue length, i.e. better utilization of the network.

VI. CONCLUSIONS AND FUTURE WORK

An adaptive decentralized strategy for active queue management of TCP flows over communication networks was proposed. The AQM strategy solves locally, at each link, an optimal control problem, minimizing residual capacity and buffer queue size. The solution of the optimal control problem is achieved via an adaptive optimization algorithm which at every times step updates the control actions in order to minimize the Hamilton-Jacobi-Bellman equation associated with the optimal control problem. Simulations results, obtained by using a fluid flow based model of the communication network, showed improvement with respect to the Random Early Detection strategy. Besides, it was shown that the performance of the proposed decentralized solution is comparable with the performance obtained with a centralized strategy, i.e. supposing the presence of a central unit that can access the state of the entire network, which highlights the feasibility of the proposed distributed approach.

Object of future work is the application of the proposed technique to a more complex discrete event description of the network, as well as to more complex network topologies. Computational complexity of implementation with respect to RED will also be investigated.

REFERENCES

- [1] V. Jacobson. Congestion avoidance and control. *Proceedings of SIGCOMM 88*, pages 314–329, 1988.
- [2] S. Floyd and V. Jacobson. Random early detection gateways for congestion avoidance. *IEEE/ACM Transactions on Networking*, 1:397–413, 1993.
- [3] F. P. Kelly, A. K. Maulloo, and D. H. K. Tan. Rate control for communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research Society*, 49:237–252, 1998.
- [4] C. V. Hollot, V. Misra, D. Towsley, and W. Gong. Analysis and design of controllers for aqm routers supporting tcp flows. *IEEE Transactions on Automatic Control*, 47:945–959, 2002.
- [5] F. P. Kelly. Mathematical modelling of the internet. In B. Engquist and W. Schmid, editors, *Mathematics Unlimited 2001 and Beyond*, pages 685–702. Springer-Verlag, Berlin, 2001.
- [6] S. H. Low, F. Paganini, and J. C. Doyle. Internet congestion control. *IEEE Control Systems Magazine*, 22:28–43, 2002.
- [7] K. Fall and K. Varadhan. *The ns manual*. available at <http://www-mash.cs.berkeley.edu/ns>, 2003.
- [8] S. Deb and R. Srikant. Global stability of congestion controllers for the internet. *IEEE Transactions on Automatic Control*, 48:1055–1060, 2003.
- [9] M. Peet and S. Lall. Nonlinear stability analysis for a class of tcp/aqm networks. *Proc. IEEE Conf. Decision Control*, pages 2309–2314, 2001.
- [10] A. Papachristodoulou, J. C. Doyle, and S. H. Low. Analysis of nonlinear delay differential equation models of tcp/aqm protocols using sums of squares. *Proc. IEEE Conf. Decision Control*, pages 14–17, 2004.
- [11] X. Fan, M. Arcak, and J. Wen. Robustness of network flow control against disturbances and time-delay. *Systems & Control Letters*, 53:13–29, 2004.
- [12] M. Peet and S. Lall. Global stability analysis of a nonlinear model of internet congestion control with delay. *IEEE Transactions on Automatic Control*, 52:553–559, 2007.
- [13] W. C. Feng, D. D. Kandlur, D. Sha, and K. G. Shin. A self-configuring red gateway. *Proceedings of IEEE INFOCOM 1999*, pages 1320–1328, 1999.
- [14] D. Lin and R. Morris. Dynamics of random early detection. *Proceedings of ACM SIGCOM 88*, pages 127–137, 1997.

(a) Residual capacity $e(t)$ (upper plot) and Queue length $b(t)$ (lower plot) at bottleneck link

(b) Sending rate $s(t)$, identical for every source (upper plot) and Probability marking $p(t)$ at bottleneck link (lower plot)

Fig. 5: Single bottleneck: $\tau_i = 40$ ms, $c_j = 9$ Mb/s. Network states under RED strategy (dotted), decentralized AO (dash-dotted), centralized AO (solid).

Table II: AO Improvements (Single Bottleneck)

CENTRALIZED	AO	BOTTLENECK
Round trip time	Capacity	Improvement wrt RED
40 ms	9 Mb/s	80.9%
200 ms	9 Mb/s	73.4 %
200 ms	18 Mb/s	72.6 %
DECENTRALIZED	AO	
Round trip time	Capacity	Improvement wrt RED
40 ms	9 Mb/s	80.7 %
200 ms	9 Mb/s	72.3 %
200 ms	18 Mb/s	67.2 %

- [15] R. J. Gibbens and F. P. Kelly. Resource pricing and the evolution of congestion control. *Automatica*, 35:1969–1985, 1999.
- [16] S. H. Low, F. Paganini, J. Wang, and J. C. Doyle. Linear stability of tcp/red and a scalable control. *Computer Networks*, 43:633–647, 2003.
- [17] H. Balakrishnan, N. Dukkipati, N. McKeown, and C. J. Tomlin. Stability analysis of explicit congestion control protocols. *IEEE Communications Letters*, 11:823–825, 2007.
- [18] M. Lestas, A. Pitsillides, P. A. Ioannou, and G. Hadjipollas. Adaptive congestion protocol: A congestion control protocol with learning capability. *Computer Networks*, 51:3773–3798, 2007.
- [19] G. Veciana, C. F. Su, and J. Walrand. Explicit rate flow control for abr services in atm networks. *IEEE/ACM Transactions on Networking*, 8:350–361, 2000.
- [20] C. Fulton, S. Li, and C. S. Lim. An abr feedback control scheme with tracking. *Proc. IEEE INFOCOM97*, 2:805–814, 1997.
- [21] P. F. Quet and H. Ozbay. On the design of aqm supporting tcp flows using robust control theory. *IEEE Transactions on Automatic Control*, 49:1031–1036, 2004.
- [22] M. Lestas, A. Pitsillides, P. A. Ioannou, and G. Hadjipollas. A new estimation scheme for the effective number of users in internet congestion control. *IEEE/ACM Transactions on Networking*, 19:1499–1512, 2011.
- [23] S. Baldi, I. Michailidis, E. B. Kosmatopoulos, A. Papachristodoulou, and P. A. Ioannou. Convex design control for practical nonlinear systems. *IEEE Trans. on Automatic Control*, 57:1692–1705, 2014.
- [24] H. Robbins and D. Siegmund. A convergence theorem for nonnegative almost supermartingales and some applications. In J. S. Rustari, editor, *Optimizing methods in statistics*, pages 233–257. New York: Academic Press, 1971.
- [25] M. D. S. Aliyu. *Nonlinear $\mathcal{H}_\infty$ Control, Hamiltonian Systems and Hamilton-Jacobi Equations*. CRC Press, Taylor & Francis Group, Boca Raton, FL, 2011.