END-TO-END RATE CONTROL FOR NETWORKS WITH RANDOM ACCESS LINKS

Clement Yuen and Peter Marbach
Department of Computer Science, University of Toronto
{clhyuen,marbach}@cs.toronto.edu

Abstract

Rate control for point-to-point wide area network has been extensively studied and is well understood by now. However, rate control for random access networks, and end-to-end congestion control for both point-to-point wide are networks and random access local area networks, has received less attention. In earlier work, we proposed and analyzed a rate control scheme for random access networks. In this paper, we propose an approach to combine this rate control scheme with existing rate control mechanism for point-to-point networks to provide end-to-end rate control.

Keywords: *Rate Control, Random Access Networks.*

1. INTRODUCTION

Rate control for wide area (point-to-point) networks has been extensively studied and is well understood by now [1–4]. Currently, the Internet uses TCP to control input rates; however TCP has some well-known drawbacks [2]. Motivated by the work of Kelly *et al* [3, 5], "price-based" rate control schemes have recently received considerable attention [6–8]. Price-based schemes have all the desirable properties of a good rate control mechanism: they allow a distributed implementation, and can be used to achieve a weighted proportional fair or approximate (arbitrarily closely) a max-min fair allocation of bandwidth. Moreover, by suitably choosing the price functions, price-based schemes can virtually eliminate packet loss and provide low-delay connections. Surprisingly, little attention has been paid to rate control for local area networks, arguably less to end-to-end rate control for networks comprising both wide area and local area networks.

In earlier work, we proposed and analyzed a rate control scheme for local area networks that support the CSMA family of random access protocols [9]. These protocols emphasize zero-coordination by allowing nodes to contend for channel access at will ("free-for-all" approach) [9]. However, if two or more nodes transmit a packet at the same time, then the packets interfere ("collide") and are not successfully received by the destination nodes. Such packets become backlogged and must be retransmitted at a later time. An important feature of the above random access network is that the network can become unstable (*i.e.* the number of backlog packets grows without bounds) if the packet generation rate exceeds the channel capacity [9]. To address this issue, we proposed a rate control mechanism to avoid an unstable operation if the unconstrained input rates exceed the channel capacity. Compared with point-to-point networks, rate control for random access networks poses additional challenges, as the aggregate input rate and the total backlog cannot be observed directly. To allow a distributed implementation, we require that the control mechanism uses only local information available at each node, namely the channel feedback information indicating whether the channel is idle, or whether a successful transmission or a collision between two or more packets has occurred. In the event of a collision we require that nodes decrease their transmission rates to avoid an increase in the backlog and additional collisions in the future. When the channel is idle, we allow nodes to increase their rates. We analyzed the above rate control scheme using the standard slotted Aloha model with an infinite set of nodes [10, 11]. For this case, we obtained the following results. We show that (1) the proposed rate control mechanism achieves a stable operation of the network in the face of excessive bandwidth demand; (2) (under appropriate assumptions) there exists a unique operating point which can be used to characterize the average throughput and packet delay; and (3) the operating point can be chosen a priori by suitably designing the rate control mechanism. Here, we discuss how this rate control scheme can be combined with price-based rate control scheme for point-to-point networks to provided end-to-end rate control.

2. RATE CONTROL IN RANDOM ACCESS NETWORKS

Below, we describe the random access network model and the rate control algorithm that we consider [9]. In addition,

CCECE 2004 - CCGEI 2004, Niagara Falls, May/mai 2004
0-7803-8253-6/04/$17.00

we outline the results that we obtained in [10].

2.1. Random Access Network Model

We model the random access network using the well-known slotted Aloha channel [9]. Consider the situation where many nodes communicate over a slotted multiaccess broadcast channel [9]. All packets have the same length and each packet requires one time unit for transmission. All nodes are synchronized such that the transmission of each packet starts at an integer time $t \geq 0$. We refer to the slot beginning at time t as slot t. If two or more nodes transmit a packet in a given time slot, then there is a collision and the receiver obtains no information about the contents or source of the transmitted packets. Packets that experience a collision become backlogged and are retransmitted at a later time. When all nodes involved in a collision immediately retransmit, then another collision would surely occur. Instead, each backlogged node retransmits a packet with a fixed probability q, $0 < q < 1$, in each successive slot until a successful transmission occurs. If just one node sends a packet in a given time slot, the packet is successfully received by the destination node(s). We assume that packets are generated at each node by application(s). Whenever a node receives a new packet from an application, it transmits the packet in the first slot after its arrival, thus risking occasional collisions but achieving a small delay if collisions are rare.

2.2. Rate Control

If the rate at which applications generate packets is left uncontrolled, then the above random access channel can get clogged with collisions. Here we consider a rate control mechanism which uses the channel output as feedback information. The intuition behind our approach is that nodes should reduce their rates when the total number of backlogged packets becomes large, thereby avoiding collisions in the future and an even further increase in the backlog. As nodes are remote and the total backlog is not globally known, we form an estimate u, $u \geq 0$, of the system backlog which we update based on the channel output. We use the estimate u as a control signal (congestion measure) to modulate the rate at which nodes are allowed to generate new packets.

For the analysis of this rate control scheme we assume an infinite set of nodes as discussed for example in [9]; it is well-known that the infinite model provides a qualitative measure of the goodness of an algorithm and it has been widely used for the stability analysis of retransmission strategies [**?**, 12–14]. New packets are thus generated according to a Poisson process where each new packet arrives to a new node [9]. Let $\lambda(u)$ denote the aggregate packet arrival rate which is modulated by a control signal u. We refer to $\lambda(u)$ as the *rate function*. We make the following assumption.

Assumption 1: The rate function $\lambda(u)$ is bounded, continuous and strictly decreasing. Furthermore, there exists a positive constant $u_{\max}$ such that $\lambda(u) = 0$ when $u \geq u_{\max}$. Note that we assume that the arrival rate vanishes when the control signal exceeds the threshold value $u_{\max}$.

Assumption 1 allows the uncontrolled rate $\lambda(0)$ to exceed the capacity of the slotted Aloha channel, leading to an unstable operation if the control signal u is not appropriately adjusted. Note that we can trivially stabilize the channel by setting the control to $u_{\max}$, but in this case the network would be idle all the time. In the following, we aim at choosing the control u to stabilize the slotted Aloha channel while not unnecessarily waste bandwidth.

Let Z_t be the random variable indicating the total number of transmission attempts in slot t. Note that Z_t includes the transmissions of new packets as well as the retransmissions of backlogged packets. Let $I[\cdot]$ be the indicator function and let $I[Z_t = 0]$, $I[Z_t = 1]$, and $I[Z_t \geq 2]$ indicate respectively an idle slot, a successful transmission, and a collision in slot t. Let u_t be the control signal that determines the rate of arrivals of new packets in slot t. We then update the control u_t by setting

$$\begin{aligned} u_{t+1} &= \big[u_t + \alpha I[Z_t = 0] + \ldots \\ &\quad + \beta I[Z_t = 1] + \gamma I[Z_t \geq 2]\big]^+, \end{aligned} \tag{1}$$

where $[x]^+$ denotes $\max\{x, 0\}, x \in \mathbb{R}$. The parameters α, β and γ are real-valued control updates corresponding to the observed channel outputs of "idle", "success" and "collision" respectively. As the aggregate rate is a decreasing function, the control signal should be increased in the event of a collision to reduce the arrival rate of new packets and to avoid an increase in the backlog and additional collisions in the future. In the event of an idle slot the control should be lowered to allow nodes to increase their rates to make more efficient use of the channel. In the event of a successful transmission, the control can be left unchanged, or lowered (raised) in the case of a more conservative (aggressive) control scheme. In line with this intuition, we choose the control parameters such that $\alpha < 0$ and $\gamma > 0$, while β can take on any real value.

2.3. Analysis and Results

In the following, we outline stability results for the above rate control scheme and results concerning the use of control parameters to achieve a desired throughput. These results can be obtained by analyzing the Markov chain representation of the system.

Let n_t be the total number of backlogged packets at the beginning of slot t. Note that we have that

$$n_{t+1} = n_t + Y_t - I[Z_t = 1], \tag{2}$$

where Y_t is a random variable indicating the number of newly arrived packets in slot t. Eq. (1) and (2) describe the dynamics of the control signal and backlog, and the system can be modeled as a Markov chain with a countable state space.

Proposition 1: The Markov chain given by Eq. (1) and (2) is positive recurrent.

Having established stability, in the next step we would like to characterize the expected throughput and delay of the system. We use the following notation. Given a state (n, u), consider the functions $d_n(n, u)$ and $d_u(n, u)$ given by

$$d_n(n,u) \triangleq \mathbf{E}\big(n_{t+1} - n_t \bigm| n_t = n, u_t = u\big),$$
$$d_u(n,u) \triangleq \mathbf{E}\big(u_{t+1} - u_t \bigm| n_t = n, u_t = u\big).$$

Definition 1: We call $(n^*, u^*) \in \mathbb{R}_+^2$ an operating point of the system if $d_n(n^*, u^*) = d_u(n^*, u^*) = 0$, We refer to $G^* = \lambda(u^*) + n^* q$ as the offered load at the operating point (n^*, u^*).

We note that the offered load G^* characterizes both the average throughput and delay at an operating point [10]. We have the following result.

Proposition 2: There always exists an operating point. Furthermore, when we have that $\lambda\big(\max\{-\alpha, -\beta\}\big) \geq e^{-1}$, then the operating point is unique.

The next lemma establishes we can set offered load at the operating point by suitably choosing the parameters α, β, and γ.

Lemma 1: Let G be a given offered load, and let $\alpha < 0$ and $\gamma > 0$ be given control parameters. For

$$\beta = \frac{\gamma}{G}(G + 1 - e^G) - \frac{\alpha}{G},$$

the offered load G^* at every operating point is such that

$$G^* = G \qquad \text{if} \quad \lambda\big(\max\{-\alpha, -\beta\}\big) \geq e^{-1},$$
$$G^* \leq G \qquad \text{if} \quad \lambda\big(\max\{-\alpha, -\beta\}\big) < e^{-1}.$$

For a detailed derivation and discussion of the above results, we refer to [10].

3. END-TO-END RATE CONTROL

In this section, we discuss how the proposed rate control scheme for random access networks can be integrated with control schemes for point-to-point wide area networks to provide end-to-end rate control.

3.1. RATE CONTROL FOR POINT-TO-POINT NETWORKS

Rate (congestion) control mechanisms for point-to-point networks have been extensively studied over the years. Here, we consider a class of rate control algorithms originally proposed by Kelly *et al* [3,4] which aims at maximizing the total user utility. Different solution approaches to this optimization problem have been studied in [6–8]. Here we focus on an approach proposed by Athuraliya and Low [15].

Consider a point-to-point network consisting of a set of links L with finite capacity c_l, $l \in L$. It is shared by a set S of sources. A source s traverses a subset $L(s) \subset L$ of links to the destination, and attains a utility $U_s(\lambda_s)$ when it transmits at rate λ_s, $\lambda_s \geq 0$. We assume that the utility U_s is increasing and strictly concave in its argument. Athuraliya and Low propose then the following algorithm to determine the transmission rate of each source.

Each link $l \in L$ updates a link congestion measure $u_{l,t}$ in time slot t based on the aggregate input rate $\lambda_{l,t}$ and the backlog $n_{l,t}$ at link l,

$$u_{l,t+1} = \big[u_{l,t} + \gamma(\alpha_l n_{l,t} + \lambda_{l,t} - c_l)\big]^+ \tag{3}$$

where $\gamma > 0$ and $\alpha_l > 0$ are small constants. Note that in point-to-point networks the input rate $\lambda_{l,t}$ and backlog $n_{l,t}$ can be observed at each link and do not have to be estimated as in our approach. Roughly, the value of u_l will increase as the backlog n_l increases. Link l marks each packet arriving in slot t that is not already marked at an upstream link, with a probability $m_{l,t}$ that is exponentially increasing in the control signal $u_{l,t}$:

$$m_{l,t} = 1 - \phi^{-u_{l,t}},$$

where $\phi > 1$ is a constant. Once a packet is marked, its mark is carried to the destination and then conveyed back to the source via acknowledgment.

The end-to-end probability that a packet of source s is marked is then given by

$$m_{s,t} = 1 - \prod_{l \in L(s)} (1 - m_{l,t}) = 1 - \phi^{-u_{s,t}},$$

where $u_{s,t} = \sum_{l \in L(s)} u_{l,t}$ is the sum of the congestion measures along the path of source s, a path congestion measure.

Source s estimates the end-to-end marking probability $m_{s,t}$ by the fraction $\hat{m}_{s,t}$ of its packets marked in slot t, and estimates its path congestion measure $u_{s,t}$ by setting

$$\hat{u}_{s,t} = -\log_\phi(1 - \hat{m}_{s,t}). \tag{4}$$

It then sets its transmission rate to $\lambda_s(\hat{u}_{s,t})$, where

$$\lambda_s(u_s) = U_s'^{-1}(u_s)$$

and $U_s'^{-1}$ is the inverse of the derivative of the utility U_s. Under the above assumptions on U_s, we have that λ_s exists and is strictly decreasing in u_s. Hence, the above algorithm says: if the path $L(s)$ is congested ($u_{s,t}$ is large), transmit at a small rate, and vice versa.

It can be shown that this scheme converges to an operating point $(u_l^*)_{l \in L}$ which has the property that

$$u_l^* \geq 0 \qquad \text{if } \lambda_l^* = c_l, \tag{5a}$$
$$u_l^* = 0 \qquad \text{if } \lambda_l^* < c_l. \tag{5b}$$

For source s, let $u_s^* = \sum_{l \in L(s)} u_l^*$ be the path congestion measure at an operating point. The rate vector at an operating point is then given by $\left(\lambda_s(u_s^*)\right)_{s \in S}$. It can be shown that this vector leads to a rate allocation which maximizes the total source utility. We note that there may be multiple operating points; however the rate allocation given by the vector $\left(\lambda_s(u_s^*)\right)_{s \in S}$ is unique (see [6, 15] for details).

3.2. Integrating Rate Control for Random Access and Point-to-Point Networks

Consider now the situation where sources (hosts) are connected to a wide area point-to-point network via local area random access networks. To achieve end-to-end rate control, we can combine the above rate control algorithm for point-to-point networks and the rate control scheme presented in this paper as follows.

Let A be the set of local area networks. A source s traverses a subset $A(s) \subset A$ of local area networks and a subset $L(s) \subset L$ of links. Let $\lambda_s(u_s)$ be the rate function of source s where

$$u_s = \sum_{a \in A(s)} u_a + \sum_{l \in L(s)} u_l,$$

and u_a and u_l are the control signals (congestion measures) associated with local area network $a \in A$ and link $l \in L$ respectively. Furthermore, suppose that the control signal u_l are updated as given by Eq. (3) and the control signal u_a according to

$$\begin{aligned} u_{a,t+1} \quad = \quad & \big[u_{a,t} + \alpha_a I[Z_{a,t} = 0] + ... \\ & \beta_a I[Z_{a,t} = 1] + \gamma_a I[Z_{a,t} \geq 2]\big]^+, \end{aligned}$$

where $Z_{a,t}$ is the total number of transmission attempts in a during slot t, and the parameters α_a, β_a, and γ_a are chosen to achieve a certain offered load G_a.

To communicate the control signals u_a and u_l to the sources, we can again use the marking scheme proposed by Athuraliya and Low [15], where in each local area network $a \in A$ all unmarked arriving packets in slot t are marked with a probability $m_{a,t}$ given by

$$m_{a,t} = 1 - \phi^{-u_{a,t}}.$$

Note that arriving packets that are sent between nodes on the same local area network are also marked with the above probability. Source s can then again estimate the end-to-end marking probability $m_{s,t}$ by the fraction $\hat{m}_{s,t}$ of its packets being marked in slot t, and estimates the control signal $u_{s,t}$ using Eq. (4).

The operating point analysis from [15] and this paper can be combined to obtain the following result for the above end-to-end rate control scheme. An operating point $\left((u_a^*)_{a \in A}, (u_l^*)_{l \in L}\right)$ exists. In addition, every operating point $\left((u_a^*)_{a \in A}, (u_l^*)_{l \in L}\right)$ has the following properties: Eq. (5) holds for all $l \in L$, and for all $a \in A$ we have

$$\begin{aligned} u_a^* \geq 0 \qquad & \text{if } G_a^* = G_a, \\ u_a^* = 0 \qquad & \text{if } G_a^* < G_a, \end{aligned}$$

where G_a^* is the offered load at the operating point. Furthermore, the rate vector at the operating point maximizes the total source utility as discussed in [6]. Again, there might be multiple operating points; however the rate allocation given by $\left(\lambda_s(u_s^*)\right)_{s \in S}$ is unique.

References

[1] V. Jacobson, "Congestion Avoidance and Control," in *Proceedings of Symposium on Communications Architectures and Protocols (SIGCOMM)*. August 1988, pp. 314–329, ACM Press.

[2] J. F. Kurose and K. W. Ross, *Computer Networking: A Top-Down Approach Featuring the Internet*, 2002. Addison Wesley, 1993.

[3] F. Kelly, A.K. Maulloo, and D.K.H. Tan, "Rate control for communication networks: shadow prices, proportional fairness and stability," *Journal of the Operational Research Society*, vol. 49, pp. 237–252, 1998.

[4] R. J. Gibbens and F. Kelly, "Resource Pricing and the Evolution of Congestion Control," *Automatica*, vol. 35, no. 12, pp. 1969–1985, December 1999.

[5] F. Kelly, "Charging and Rate Control for Elastic Traffic," *European Transactions on Telecommunications*, vol. 8, pp. 33–37, 1997.

[6] S. H. Low and D. E. Lapsley, "Optimization Flow Control, I: Basic Algorithm and Convergence," *IEEE/ACM Transactions on Networking*, vol. 7, no. 6, pp. 861–874, December 1999.

[7] S. J. Golestani and S. Bhattacharyya, "A class of end-to-end congestion control algorithms for the Internet," in *Proceedings of the Sixth International Conference on Network Protocols*, October 1998, pp. 137–150.

[8] S. Kunniyur and R. Srikant, "End-to-end Congestion Control Schemes: Utility Functions, Random Losses and ECN Marks," in *Proceedings of the Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM 2000)*, March 2000, pp. 1323–1332, vol. 3.

[9] D. Bertsekas and R. Gallager, *Data Networks*, Prentice-Hall, Inc., second edition, 1992.

[10] C. Yuen, "Rate Control with Pricing in Contention-Based Access Networks," Master thesis, University of Toronto, 2002.

[11] C. Yuen and P. Marbach, "Service Differentiation in Random Access Networks," Technical Report CSRG-472, University of Toronto, July 2003.

[12] B. and T. van Loon, "Decentralized Dynamic Control of a Multiaccess Broadcast Channel," *IEEE Transactions on Automatic Control*, vol. AC-27, no. 3, pp. 559–569, June 1982.

[13] F. Kelly, "Stochastic Models of Computer Communication Systems," *Journal of the Royal Statistical Society*, vol. B 47, no. 3, pp. 379–395, 1985.

[14] Loren P. Clare, "Control procedures for slotted Aloha systems that achieve stability," in *Proceedings of the ACM SIGCOMM on communications architecture & protocols*, Stowe, VT, August 1986, pp. 302–309.

[15] S. Athuraliya and S. H. Low, "Optimization Flow Control, II: Implementation," May 2000, Submitted for publication.