A Robust Flow Control Framework for Heterogenous Network Access

Tansu Alpcan[1], Jatinder Pal Singh[1], and Tamer Başar[2]

[1] Deutsche Telekom Laboratories, Technische Universität Berlin
Ernst-Reuter Platz 7, 10587 Germany. Email: *(tansu.alpcan, jatinder.singh)@telekom.de*
[2] Coordinated Science Lab., University of Illinois
1308 West Main St. Urbana, IL 61801, USA. Email: *tbasar@control.csl.uiuc.edu*

***Abstract*— We investigate a novel robust flow control framework for heterogeneous network access by devices with multihoming capabilities. Towards this end, we develop an H^∞-optimal control formulation for allocating rates to devices on multiple access networks with heterogeneous time-varying characteristics. H^∞ analysis and design allow for the coupling between different devices to be relaxed by treating the dynamics for each device as independent of the others. Thus, the distributed end-to-end rate control scheme proposed in this work relies on minimum information and achieves fair and robust rate allocation for the devices. An efficient utilization of the access networks is established through an equilibrium analysis in the static case. We perform measurement tests to collect traces of the available bandwidth on various WLANs and Ethernet. Through simulations, our approach is compared with AIMD and LQG schemes. In addition, the efficiency, fairness, and robustness of the H^∞-optimal rate controller developed are demonstrated via simulations using the measured real world network characteristics.**

I. Introduction

Contemporary networks are heterogeneous in their attributes such as the supporting infrastructure, protocols, and offered data rates. The multitude and variety of existing and emerging wireless and wired networking technologies continue to be the driving force towards convergence of networks. It is commonplace today to have electronic devices with multiple networking capabilities. Personal computing devices, e.g., laptops, PDAs, smartphones, are typically equipped with several access systems ranging from different types of IEEE 802.11 wireless local area networks (WLAN) to Ethernet, GPRS, and UMTS.

On end-user devices a variety of applications emerge with different bandwidth requirements for multimedia access, gaming, and collaboration. Our objective is efficient utilization of multiple networks by devices via rate control and optimal assignment of traffic flows to available networks. The functionality we envision can be described in a hypothetical scenario. Imagine a user in a corporate setting participating in a video conference call via her device having both Ethernet and WLAN (say IEEE 802.11g) connectivity. While engaged in the conference proceedings, the user is uploading content on a remote server for the participants to access, and at the same time needs to retrieve some files from the server. Several traffic flows are hence created by the device which dynamically monitors the networks at its disposal. The device then routes the flows via these networks and dynamically reassigns them to different networks based on the varying network characteristics like available bit rate (ABR) and delay.

While the distribution of traffic flows amongst different networks can enable better network utilization than single network use at a time, the variation in network characteristics like ABR and delay makes the problem of flow control and assignment challenging. This and similar problems have been explored from different perspectives. A game theoretic framework for bandwidth allocation for elastic services in networks with fixed capacities has been addressed in [1]–[3]. Packet scheduling for utilization of multiple networks has been investigated in [4]. Flow scheduling for collaborative Internet access in residential areas via multihomed client devices is discussed in [5]. A solution for addressing the handoff, network selection, and autonomic computation for integration of heterogeneous wireless networks has been presented in [6]. The work, however, does not address efficient simultaneous use of heterogeneous networks and does not consider wireline settings. In [7], the authors have explored design of a network comprised of wide area and local area technologies where user devices select among the two technologies in a greedy fashion so as to maximize a utility function based on wireless link quality, network congestion, etc. Recently, a cost price mechanism that enables a mobile device to split its traffic amongst several IEEE 802.11 access points based on throughput obtained and price charged, was proposed in [8]. In [9], an optimal rate control scheme has been investigated where the queue size at a bottleneck link is optimized using an H^∞ control formulation.

In this paper, we address the problem of optimal rate control and assignment of flows on a device onto multiple networks with randomly varying characteristics by developing principles for an envisioned *middleware* functionality. The main motivation for our approach is the practically observed varying characteristics of networks. We consider a setting where multiple devices hosting a variety of flows with different requirements have access to several different access networks. Each device has to decide on its flow rate on each network

[1]Research supported by Deutsche Telekom AG.

which it determines by measuring the available bandwidth. Here, unlike our recent work [10], the devices do not capture the characteristics of the network, which vary randomly over time within a Markov model, but use a linear state-space system to keep track of current and past observations on available bandwidth. We consider a worst-case formulation and let the devices update their flow rates using various formulations of H^∞-optimal control. The use of H^∞ analysis allows to treat the dynamics for each device as independent of the others and the variations in the available bandwidth are modeled as unknown disturbances. We study the control schemes proposed both analytically and numerically via simulations in Matlab.

The rest of the paper is organized as follows: We describe the network access problem and our measurement results in Section II. The network model is presented in Section III. The analytical framework for robust flow control is discussed in Section IV, which also includes an equilibrium analysis and an illustrative example. Section V generalizes the formulation in Subsection IV-A. We describe the simulations and results obtained therefrom in Section VI. The paper concludes with remarks and a discussion of future research directions in Section VII.

II. Heterogenous Network Access and Measurements

We describe the underlying architecture for routing and rate control of flows originating from applications running on a device via access networks that are available to the device. Essential to the design is a *middleware* that runs a lightweight tool to estimate the uplink and downlink bandwidths and delay for the flows on different networks. Applications running on the device consult the middleware for routing of flows to and from the corresponding hosts in the Internet via different networks. For the uplink flows, network assignment is done by assigning flows generated by the applications on the device to suitable networks by consulting the middleware. For the downlink flows, suitable network assignment can be done via handshake mechanisms with the sending host in the Internet. This handshaking would involve the sending host assisting the device in assigning downlink flows to the network suggested by the middleware.

An interesting question that arises in this setting is the characterization of the access networks in terms of bandwidth and delay. We conduct network measurements in a real world setting and will use the traces to simulate (Section VII) the flow control schemes introduced in the subsequent sections. We monitor the ABR and round trip time (RTT) on different networks including Ethernet and IEEE 802.11b and IEEE 802.11g WLANs. The tests are conducted between hosts in Deutsche Telekom Laboratories (T-Labs) in Berlin to three destinations - Stanford University, Technical University of Munich (TU Munich), and the Technical University of Berlin (TU Berlin) - respectively representing long, mid, and close distance destinations [10]. We surveyed several publicly available tools including Pathrate, Nettest, CapProbe and chose Abing [11] for measurement of ABR and RTT.

Table I shows the average ABR and RTT and their standard deviations to the TU Berlin destination for different networks for 2 hour traces. We observe that all the networks display noticeable variation in ABRs. For instance the ABR on 802.11g can be as high as 24 Mbps and can drop down to as low as 6 Mbps. These random fluctuations in the ABR observed experimentally clearly indicate the need for a robust flow control approach.

TABLE I
Available bit rate and RTT from T-Labs to TU Berlin

		ABR(Mbps)	RTT(ms)
Ethernet	Avg.	71.8	5.2
	Std. Dev.	13.0	0.04
802.11g	Avg.	14.3	7.8
	Std. Dev.	3.6	0.4
802.11b	Avg.	4.5	10.7
	Std. Dev.	0.5	0.6

III. Network Model

In this section we present an analytical model of the heterogenous network (access) environment. We consider a set of access networks $\mathcal{I} := \{1, 2, \ldots, I\}$ simultaneously available to multiple devices. Let us define the set of such devices sharing these networks as $\mathcal{D} := \{1, 2, \ldots, D\}$. The assignment and control of flows originating from these devices on these access networks constitute the underlying resource allocation problem we investigate in this paper. Let the nonnegative flow rate a device $d \in \mathcal{D}$ assigns to each available network $i \in \mathcal{I}$ be $\mathbf{r}_d := [r_d^{(1)}, r_d^{(2)}, \ldots, r_d^{(I)}]$. Consequently, the total flow rate on network i is $\bar{r}^{(i)} = \sum_{d=1}^{D} r_d^{(i)}$. An important property of the access networks investigated in this paper is the high variability of the network capacity $C^{(i)}(t)$, where t denotes time. The available bandwidth $B^{(i)}$ on a network is then given by $B^{(i)}(t) := C^{(i)}(t) - \sum_{d=1}^{D} r_d^{(i)}(t)$. A device can estimate via various online measurement tools [11] the quantity $w^{(i)} := \phi\left(B^{(i)}(t)\right)$, where the function $\phi(\cdot)$ is approximately proportional to its argument, i.e. the available bandwidth $B^{(i)}$.

In our recent study [10], we have investigated a Markov chain based framework for modeling the access network properties and addressed mainly the problem of discrete flows assignment. In this paper we shift our focus to explicit flow control. Towards this end, we introduce and investigate in the next section a linear system formulation and an H^∞ controller that optimizes the network usage.

IV. Robust Flow Control

Most of the access networks available to a device at a given time are wireless ones. The characteristics of a (wireless) network $i \in \mathcal{I}$, and hence, its available bandwidth $B^{(i)}$ fluctuates randomly due to fading effects in the case of wireless networks as well as background traffic. Following a different path than the one in [10] we do not attempt to model this quantity but take a function of it $w^{(i)}(B^{(i)})$ simply as an input to devices. We do not make any assumptions on the nature of this function which captures the random variations in available bandwidth due to channel state and other factors.

We next define a system from the perspective of a device $d \in \mathcal{D}$ which keeps track of the available bandwidth of a single access network. The system state $x_d^{(i)}$ reflects from the perspective of device d roughly the bandwidth availability on network i. In order to simplify the analysis we focus in this section on the single network case and drop the superscript i for notational convenience. Then, the system equation for device d is

$$\dot{x}_d = a\, x_d + b\, u_d + w, \tag{1}$$

where u_d represents the control action of the device. The parameters $a < 0$ and $b < 0$ adjust the memory horizon (the smaller a the longer the memory) and the "expected" effectiveness of control actions, respectively, on the system state x_d. The device d bases its control actions on its state which not only takes as input the current available bandwidth but also accumulates the past ones to some extent. It is also possible to interpret the system (1) as a low pass filter with input w and output x.

Let us introduce a rate update scheme which is approximately proportional to the control actions:

$$\dot{r}_d = -\phi r_d + u_d, \tag{2}$$

where $\phi > 0$ is sufficiently small. Although this rate update scheme seems disconnected from the system in (1) it is not the case as we show in Section IV-B. As a result of w being a function of the available bandwidth B, which in turn is a function of the aggregate user rates, the systems (1) and (2) are connected via a feedback loop.

For simplicity the coefficient of u_d is chosen to be one in (2). Since a rate update of a device will have the inverse effect on the available bandwidth, the parameter b in (1) has to be then negative. We additionally note that we resort here to a "bandwidth probing scheme" in a sense similar to additive-increase multiplicative-decrease (AIMD) of the well-known transfer control protocol (TCP). On the other hand, our scheme's main parameters follow from an optimization problem which will be defined next.

We now address the question of how to calculate the rate control action u_d of a device. To be able to do so, we first need to formulate our objectives based on the full bandwidth-utilization criterion. We make the following observations on the system (1) and (2): First, the input w is zero if all devices fully use the available bandwidth on the access network. Second, the system is stable, i.e., the state converges to zero unless w and u_d are nonzero. Third, the rate change is a function of control actions u_d. Finally, the control actions u_d have a direct effect on the state x_d. Thus, we can formulate the objectives of the optimal controller as minimizing (squares of) w, the state x, and the control actions u. These objectives ensure that the input or "disturbance" to the system is rejected (maximum capacity usage) while preventing excessive rate fluctuations leading to instabilities and jitter.

A. H^∞-optimal control

In order to achieve the objectives defined above on system (1) without explicity making any assumptions on w we utilize H^∞ optimal control theory. It provides a powerful framework that allows for a worst-case analysis of disturbance attenuation problems. H^∞ analysis allows for the coupling between different devices to be ignored. Furthermore, the dynamics for each device are treated as independent of the others and driven by unknown disturbances. By viewing the disturbance (here the available bandwidth) as an intelligent maximizing opponent in a dynamic zero-sum game who plays with knowledge of the minimizer's control action, we evaluate the system under the worst possible conditions (in terms of capacity usage). We then determine the control action that will minimize costs or achieve the objectives defined under these worst circumstances [12].

The system described can be classified as continuous-time with perfect state measurements due to the state x_d being an internal variable of device d. We conduct an H^∞-optimal control analysis and design taking this into account. Let us first introduce the *controlled output*, $z_d(t)$, as a two dimensional vector:

$$z_d(t) := [h x_d(t) \;\; g u_d(t)]^T, \tag{3}$$

where g and h are positive parameters. The cost of a device that captures the objectives defined and for the purpose of H^∞ analysis is the ratio of the L^2-norm of z_d to that of w:

$$L_d(x_d, u_d, w) = \frac{\|z_d\|}{\|w\|}, \tag{4}$$

where $\|z_d\|^2 := \int_0^\infty |z_d|^2 \, d\tau$, and a similar definition applies to $\|w\|^2$. Although being a ratio, we will refer to L_d as the (device) cost in the rest of the analysis. It captures the proportional changes in z_d due to changes in w. If $\|w\|$ is very large, the cost L_d should be low even if $\|z_d\|$ is large as well. A large $\|z_d\|$ indicates that the state $|x_d|$ and the control $|u_d|$ have high values reflecting and reacting to the situation, respectively. However, they should not grow unbounded, which is ensured by a low cost, L_d. For the rest of the analysis, we will drop the subscript d for ease of notation.

H^∞-optimal control theory guarantees that a performance factor will be met. This factor γ, also known as the H^∞ norm, can be thought of as the worst possible value for the cost L. It is bounded below by

$$\gamma^* := \inf_u \sup_w L(u, w), \tag{5}$$

which is the lowest possible value for the parameter γ. It can also be interpreted as the optimal performance level in this H^∞ context. Interestingly, we assume here that the available bandwidth is "controlled" by a maximizing player (we call as Murphy) who plays second in this formulation knowing the control applied by the minimizing player or device. This formulation as well as the order of play ensures that we are indeed analyzing the worst case scenario.

In order to solve for the optimal controller $\mu(x)$, a corresponding differential game is defined, which is parameterized by γ,

$$J_\gamma(u,w) = \|z\|^2 - \gamma^2\|w\|^2. \tag{6}$$

The maximizing player (Murphy) tries to maximize this cost function while the objective of a device is to minimize it. The optimal control action $u = \mu_\gamma(x)$ can be determined from this differential game formulation for any $\gamma > \gamma^*$.

This controller is expressed in terms of a relevant solution, σ_γ, of a related game algebraic Riccati equation (GARE) [12]:

$$2a\sigma - \left(\frac{b^2}{g^2} - \frac{1}{\gamma^2}\right)\sigma^2 + h^2 = 0 \tag{7}$$

By the general theory [12], the relevant solution of the GARE is the "minimal" one among its multiple nonnegative-definite solutions. However, in this case, since the GARE is scalar, and the system is open-loop stable (that is, $a < 0$), the GARE (which is a quadratic equation) admits a unique positive solution for all $\gamma > \gamma^*$, and the value of γ^* can be computed explicitly in terms of the other parameters. Solving for the roots of (7), we have:

$$\sigma_\gamma = \frac{-a \pm \sqrt{a^2 - \lambda h^2}}{\lambda}$$

where

$$\lambda := \frac{1}{\gamma^2} - \frac{b^2}{g^2}.$$

The parameter λ could be both positive and negative, depending on the value of γ, but for γ close in value to γ^* it will be positive. Further, γ^* is the smallest value of γ for which the GARE has a real solution. Hence,

$$\gamma^* = \left[\sqrt{\frac{a^2}{h^2} + \frac{b^2}{g^2}}\right]^{-1}$$

Finally, a controller that guarantees a given performance bound $\gamma > \gamma^*$ is:

$$\mu_\gamma(x) = -\left(\frac{b}{g^2}\,\sigma_\gamma\right)\,x. \tag{8}$$

This is a linear feedback controller operating on the device system state x, where the gain can be calculated offline using only the linear quadratic system model and for the given system and cost parameters.

B. Equilibrium and Stability Analysis for Static Capacity

The H^∞-optimal rate controller which has been derived in the previous section has the general form $u = \theta x$, where θ is a positive constant. We conduct an equilibrium and stability analysis of the system (1) and (2) under this general class of linear feedback controllers for a single network of fixed capacity C and accessible by D devices. The results obtained also apply to the specific case of the H^∞ controller.

By ignoring the noise in the system, we make the simplifying assumption of $w := C - \sum_{i=1}^{D} r_i$ and let $d = 1$. Then,

$$\begin{aligned} \dot{x}_i &= a\,x_i + b\,\theta x_i + C - \textstyle\sum_{k=1}^{D} r_k \\ \dot{r}_i &= -\phi r_i + \theta x_i, \quad i = 1,\ldots,D. \end{aligned} \tag{9}$$

At the equilibrium -which we will show below exists, is unique, and is asymptotically stable- we have $\dot{x}_i = \dot{r}_i = 0\ \forall i$. Solving for equilibrium values of x_i and r_i for all i, denoted by x_i^* and r_i^*, respectively, it is easy to obtain

$$r_i^* = \frac{C\theta}{\theta D - (a + b\theta)\phi}$$

and

$$x_i^* = \frac{C\phi}{\theta D - (a + b\theta)\phi},$$

which are unique, under the negativity of a and b and positivity of θ, as long as $\phi > 0$. Now, we note that as $\phi \to 0^+$, we have $\sum_i r_i \to C$. Thus, as ϕ approaches to zero from the positive side, linear feedback controllers of the form $u = \theta x$, where $\theta > 0$ (including H^∞-optimal controllers) ensure maximum network usage when the capacity C is fixed and there is no noise .

We now show that the linear system (9) is stable and asymptotically converges to the equilibrium point whenever $\phi > 0$. Toward this end, let us sum the rates r_i in (9) to obtain

$$\begin{aligned} \dot{x}_i &= -\mu x_i - \bar{r} + C, \quad i = 1,\ldots,D \\ \dot{\bar{r}} &= -\phi\bar{r} + \theta \textstyle\sum_{i=1}^{D} x_i\,, \end{aligned} \tag{10}$$

where $\bar{r} := \sum_{i=1}^{D} r_i$ and $\mu := -(a + b\theta) > 0$. We can rewrite (10) in the matrix form as

$$\dot{y} = Fy + [C \cdots C\ 0]^T,$$

where $y := [x_1\ \cdots\ x_D\ \bar{r}]^T$. Then, it is straightforward to show that the characteristic function of the $(D+1)$ dimensional square matrix F has the form

$$\det(sI - F) = (s + \mu)^{D-1}\left[s^2 + (\mu + \phi)s + \mu\phi + D\,\theta\right] = 0.$$

Notice that, F has $D - 1$ repeated negative eigenvalues at $s = -\mu$ and two additional eigenvalues at

$$s_{1,2} = -\frac{1}{2}(\mu + \phi) \pm \frac{1}{2}\sqrt{(\mu - \phi)^2 - 4D\,\theta}.$$

If $(\mu - \phi)^2 < 4D\,\theta$, then both of these eigenvalues are imaginary with negative real parts. Otherwise, we have $\mu + \phi > |\mu - \phi|$ and both eigenvalues are negative and real. Therefore, all eigenvalues of F always have negative real parts and the linear system (10) is stable.

It immediately follows that x_i is always finite and converges to the equilibrium, and from the second equation of (9), r_i has to be finite and converges for all i. We thus establish that the original system (9) is stable.

C. Alternative Controller Formulations

Having established and analyzed the H$^\infty$-optimal controller for the system at hand, we study variations of it and other formulations. One possible formulation is the well-known linear quadratic Gaussian (LQG) problem where the input w is modeled as a Gaussian noise. Although this assumption probably does not hold for the problem at hand we use the LQG as a comparsion case. It can be obtained here simply as the limit of the H$^\infty$ problem as $\gamma \to \infty$. We use as the $\|z\|^2$, the expected value of $\int_0^\infty |z|^2 \, dt$, which we again denote by $\|z\|^2$ by a slight abuse of notation, and the problem is one of minimization of $\|z\|^2$.

As an alternative, we define a simple AIMD controller as another comparison scheme:

$$\dot{r}_d = \begin{cases} \alpha & , \text{if } w > 0 \\ -\beta r & , \text{if } w < 0 \end{cases}, \tag{11}$$

where α and β are positive parameters.

D. Illustrative Example

We illustrate the H$^\infty$-optimal controller with an example. The cost and system parameters are chosen simply as $a = -1$, $b = -1$, $g = 1$, $h = 1$. Then, $\gamma^* = \sqrt{2}/2 \approx 0.707$. If we choose $\gamma = \gamma^*$, then the unique positive solution of the GARE is $\bar{\sigma}_{\gamma^*} = 1$ which leads to the simple feedback controller $\mu_{\gamma^*}(x) = x$, which is the optimal H$^\infty$ controller.

We compare this result analytically with the LQG formulation where $\gamma \to \infty$. Then, the GARE, $\sigma^2 + 2\sigma - 1 = 0$, simply yields the unique positive solution $\tilde{\sigma}_\gamma \approx \sqrt{2} - 1$, which leads to $\tilde{\mu}_\gamma(x) \approx (\sqrt{2} - 1)x$. We observe that despite the same cost structure, the optimal H$^\infty$ controller is more "aggressive" in order to ensure an upper bound on the cost L regardless of w. On the other hand, the LQG controller has a lower feedback gain possibly due to the inherent Gaussian noise assumption on w.

V. H$^\infty$-OPTIMAL CONTROL FOR MULTIPLE NETWORKS

In Section IV we have provided the analysis and controller design for a single access network shared by multiple devices to increase readability and focus on core concepts by keeping the notation simple. We now provide the H$^\infty$-optimal control formulation for the general case of multiple access networks for a single device $d \in \mathcal{D}$ and drop the subscript d again for ease of notation.

Let us define $\mathbf{x} := [x^{(i)}]$, $\mathbf{r} := [r^{(i)}]$, and $\mathbf{u} := [u^{(i)}]$ for all $i \in \mathcal{I}$. Then, the counterpart of the system (1) and (2) is given by

$$\begin{aligned} \dot{\mathbf{x}} &= A\,\mathbf{x} + B\,\mathbf{u} + D\,\mathbf{w} \\ \dot{\mathbf{r}} &= -\Phi\,\mathbf{r} + \mathbf{u}, \end{aligned} \tag{12}$$

where $\mathbf{w} := [w^{(i)}] \;\; \forall i$. Here, the matrices A, B, and Φ are obtained simply by multiplying the identity matrix by a, b, and ϕ, respectively.

The counterpart of the *controlled output* in (3) is

$$\mathbf{z}(t) := H\mathbf{x}(t) + G\mathbf{u}(t), \tag{13}$$

where we assume that $G^T G$ is positive definite, and that no cost is placed on the product of control actions and states: $H^T G = 0$. The matrix H represents a cost on variation from zero state, i.e. full capacity usage. Let us in addition define $Q := H^T H$.

We next define the cost

$$L(\mathbf{x}, \mathbf{u}, \mathbf{w}) = \frac{\|\mathbf{z}\|}{\|\mathbf{w}\|}, \tag{14}$$

where $\|\mathbf{z}\|^2 := \int_0^\infty |\mathbf{z}(t)|^2 dt$, and the corresponding differential game parameterized by γ,

$$J_\gamma(\mathbf{u}, \mathbf{w}) = \|\mathbf{z}\|^2 - \gamma^2 \|\mathbf{w}\|^2 \tag{15}$$

as in Section IV-A. Here γ is larger than γ^*, where γ^* is defined as in (5).

The corresponding GARE

$$A^T Z + ZA - Z(B(G^T G)^{-1} B^T - \gamma^{-2} D D^T)Z + Q = 0, \tag{16}$$

admits a unique minimal nonnegative definite solution $\bar{Z}_\gamma$, for $\gamma > \gamma^*$, if (A, B) is stabilizable and (A, H) is detectable [12]. Thus, we obtain the H$^\infty$-optimal linear feedback controller for the multiple network case:

$$\mu_\gamma(\mathbf{x}) = -(G^T G)^{-1} B^T \bar{Z}_\gamma \mathbf{x}, \tag{17}$$

for each $\gamma > \gamma^*$, which is also stabilizing.

VI. SIMULATIONS

We simulate the H$^\infty$-optimal controller in a scenario where 20 devices share three different network interfaces with varying available bandwidths obtained from real world measurements as discussed in Section II. The parameters are $a = -1$, $b = -1$, $g = 1$ and chosen to be the same over all three networks resulting in G to be the identity matrix and A and B to be the negative identity matrices. For the last 15 devices, the Q matrix is chosen as the identity matrix while for the first 5 devices it is

$$Q = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 4 & 0 \\ 0 & 0 & 1 \end{bmatrix},$$

which indicates a preference for network 2 due to, for example, favorable delay characteristics and nature of applications running on these devices. Hence, the controllers $u1$ and $u2$ for the first 5 and last 15 devices, respectively, are

$$u1 = \begin{bmatrix} 0.5 & 0 & 0 \\ 0 & 3.9 & 0 \\ 0 & 0 & 0.5 \end{bmatrix} \cdot \begin{bmatrix} x^{(1)} \\ x^{(2)} \\ x^{(3)} \end{bmatrix},$$

$$u2 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \cdot \begin{bmatrix} x^{(1)} \\ x^{(2)} \\ x^{(3)} \end{bmatrix}.$$

The corresponding values of γ^* are calculated as 0.895 and 0.707, respectively.

The resulting aggregate flow rates and capacity of each network are depicted in Figures 1, 2, and 3. The average capacity usage on the networks are 88.2%, 89.2%, and 89.9%, respectively. The corresponding individual flows of devices on each network are shown in Figures 4, 5, and 6. As expected, the 5 devices with a preference for network 2 receive a higher share of bandwidth on it.

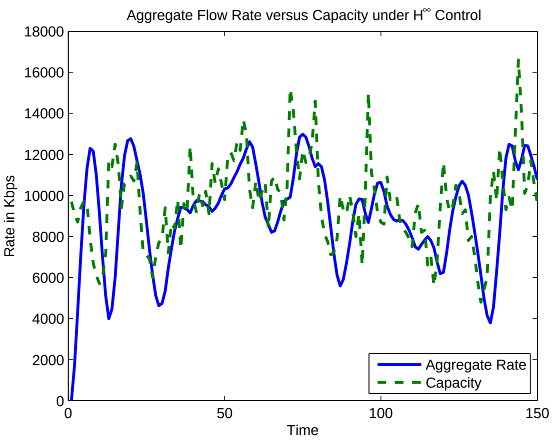

Fig. 1. The aggregate flow rate and the available capacity on network 1 under H^{∞}-optimal control.

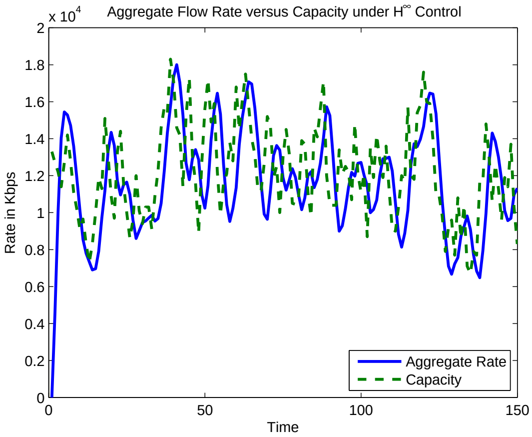

Fig. 2. The aggregate flow rate and the available capacity on network 2 under H^{∞}-optimal control.

We next compare the H^{∞}-optimal controller with the AIMD scheme in (11) with parameters $\alpha = 10$ and $\beta = 0.75$ for 20 symmetric devices on network 1. The results are depicted in Figure 7. We observe that the average capacity usage under the AIMD controller is only 74.6% and the H^{∞}-optimal controller outperforms AIMD in this aspect. In addition, the flow rates fluctuate less under the H^{∞} scheme despite a careful choice of AIMD parameters.

As a second comparison, we simulate the LQG controller discussed in IV-D within the same environment. As shown in Figure 8 the LQG performs better than the AIMD but worse than the H^{∞}-optimal controller with an average capacity usage of 85.6%. We observe that the aggregate flow rate does not follow the capacity as closely as it was the case with the

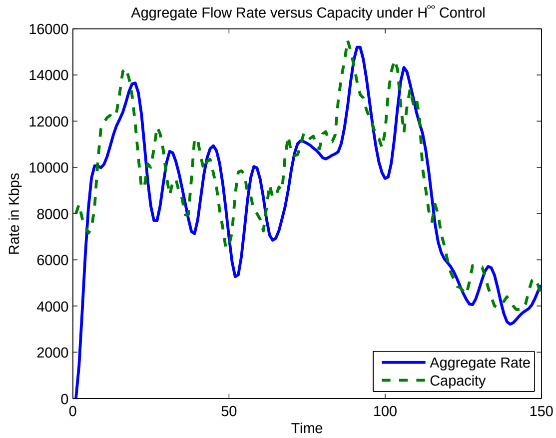

Fig. 3. The aggregate flow rate and the available capacity on network 3 under H^{∞}-optimal control.

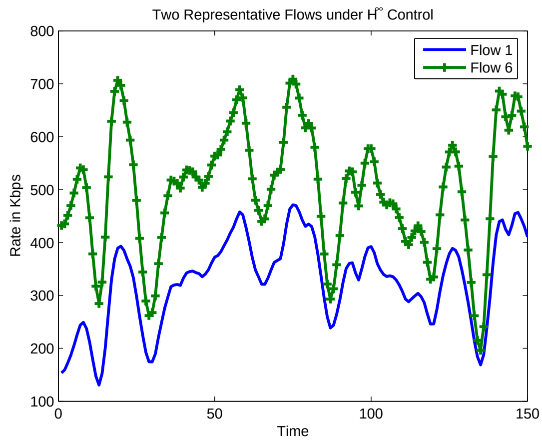

Fig. 4. The rates of representative flows 1 and 6 on network 1 under H^{∞}-optimal control.

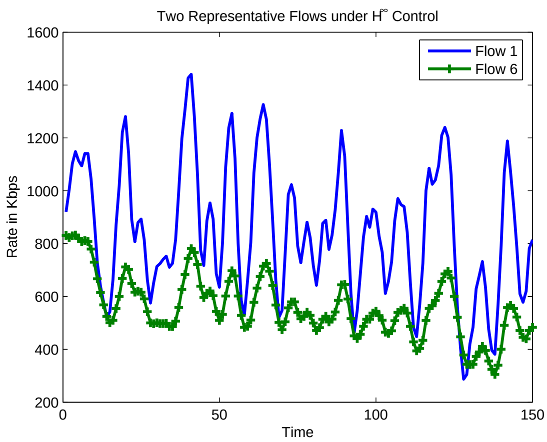

Fig. 5. The rates of representative flows 1 and 6 on network 2 under H^{∞}-optimal control.

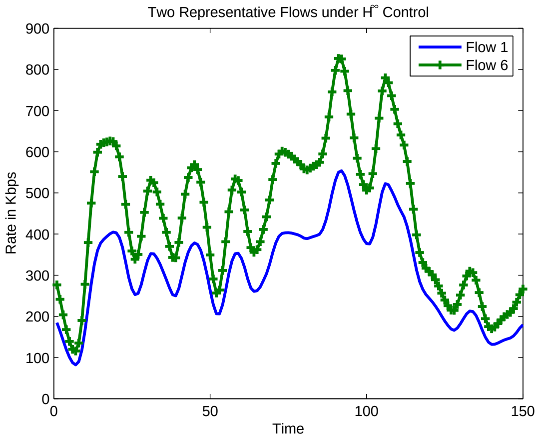

Fig. 6. The rates of representative flows 1 and 6 on network 3 under H$^\infty$-optimal control.

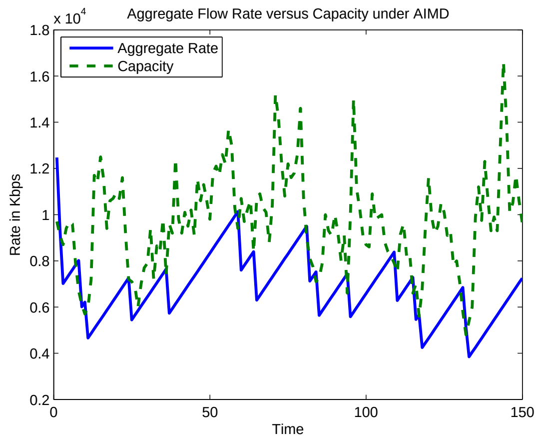

Fig. 7. The aggregate flow rate and the available capacity on network 1 under AIMD control.

H$^\infty$-optimal control. It is also important to note that the LQG scheme does not provide a minimum performance guarantee on the cost L whereas the H$^\infty$ control ensures one.

Subsequently, we investigate the fairness properties of our approach by simulating 20 devices with random initial flow rates on network 1. We observe in Figure 9 that within a short time each flow converges to an equal share of the available bandwidth or capacity. Previous simulation results in Figures 4, 5, and 6 also show that the devices which obtain more bandwidth on network 2 get less on the other two networks further indicating the fairness of our approach.

Finally, we study the robustness of H$^\infty$-optimal controller with respect to variations in the number of devices. In this simulation, we abruptly increase the number of devices accessing the network from 20 to 40 at time step $t = 100$ and deactivate them again at $t = 200$. We observe in Figures 10 and 11 that our algorithm successfully responds to these changes with a high speed of convergence.

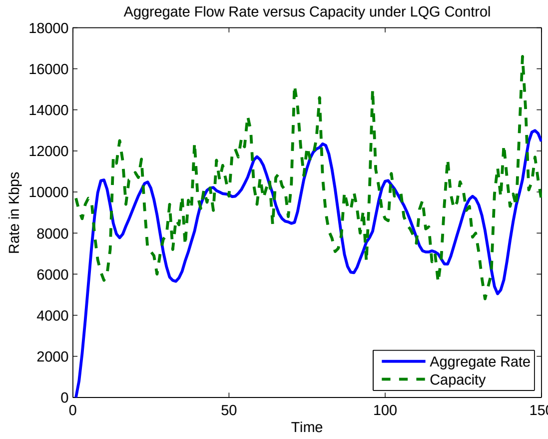

Fig. 8. The aggregate flow rate and the available capacity on network 1 under LQG control.

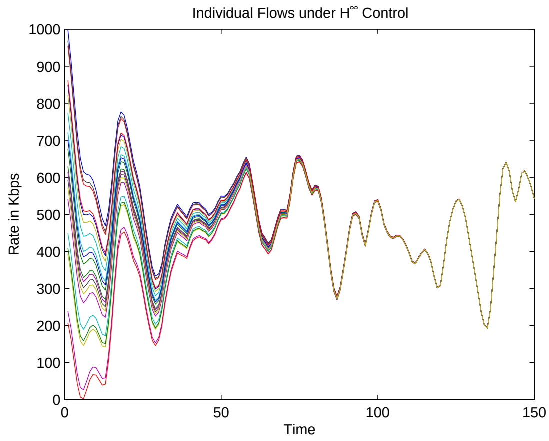

Fig. 9. The rates of individual flows with random starting points on network 1 under H$^\infty$-optimal control.

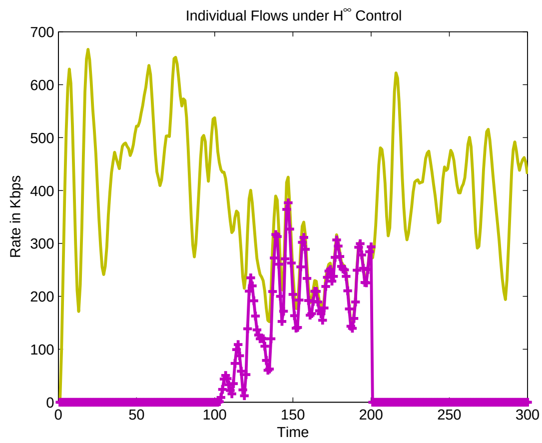

Fig. 10. The flow rates on network 1 under H$^\infty$-optimal control and varying number of active devices.

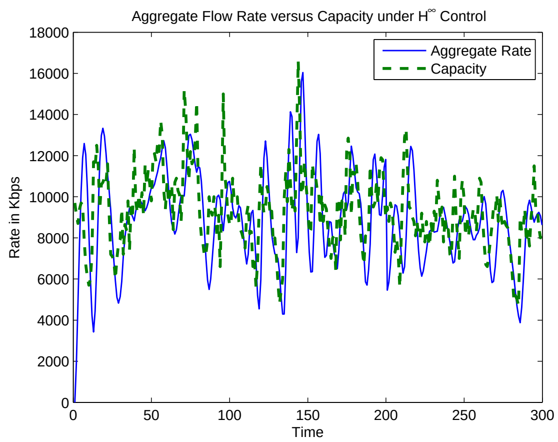

Fig. 11. The aggregate flow rate and the available capacity on network 1 under H$^{\infty}$-optimal control and varying number of active devices.

VII. Conclusion

We have presented a robust flow control approach based on H$^{\infty}$-optimal control theory for the purpose of efficient utilization of multiple heterogenous networks and a fair bandwidth allocation to devices accessing them. Bandwidth and delay measurements of different network types in a real world setting have indicated random fluctuations of these quantities and justified the necessity of a robust rate control scheme. We have modeled the system from a device's perspective and derived a minimum information rate control scheme using optimal control actions obtained through H$^{\infty}$ analysis and design. By reformulating the rate control problem as one of disturbance rejection, we have utilized H$^{\infty}$ control theory without making any restrictive assumptions on the random nature of network characteristics.

An efficient utilization of the access networks under our algorithm has been established through an equilibrium analysis in the static case. We have considered a LQG (as a variation of the H$^{\infty}$) control scheme as well as a simple AIMD algorithm for comparison purposes. The efficiency, fairness, and robustness properties of the H$^{\infty}$-optimal rate controller developed have been demonstrated via simulations using the measured real world network characteristics.

The promising results obtained are motivating for future research. One immediate direction for extension is to extend the H$^{\infty}$ analysis to the noisy measurement case where the devices estimate the available capacity on the network with some errors. Another interesting direction is the study of adaptive parameter update schemes for the linear system model of devices. Yet another extension would be to take the x_i and r_i dynamics to be device dependent, that is for the parameters a, b, ϕ and θ to be indexed by i. The equilibrium values of x_i and r_i can also be readily computed in that case, but the proof of asymptotic stability seems to be fairly involved, though tractable. Finally, we note that the robust flow control scheme can also be applied to wired networks.

References

[1] H. Yaiche, R. Mazumdar, and C. Rosenburg, "A game theoretic framework for bandwidth allocation and pricing in broadband networks," in *IEEE/ACM Transaction on Networking*, vol. 8, no. 5, Oct. 2000, pp. 667–678.

[2] T. Alpcan and T. Başar, "A utility-based congestion control scheme for Internet-style networks with delay," *IEEE Transactions on Networking*, vol. 13, no. 6, pp. 1261–1274, December 2005.

[3] ——, "Global stability analysis of an end-to-end congestion control scheme for general topology networks with delay," in *Proc. of the 42nd IEEE Conference on Decision and Control*, Maui, HI, December 2003, pp. 1092 – 1097.

[4] K. Chebrolu and R. Rao, "Communication using multiple wireless interfaces," in *Proc. IEEE Wireless Communications and Networking Conference (WCNC 2002)*, vol. 1, 2002, pp. 327–331.

[5] N. Thompson, G. He, and H. Luo, "Flow scheduling for end-host multihoming," in *Proc. IEEE INFOCOM*, 2006.

[6] P. Vidales, J. Baliosion, J. Serrat, G. Mapp, F. Stejano, and A. Hopper, "Autonomic sytem for mobility support in 4G networks," in *IEEE Journal on Selcted Areas in Communications'05*, vol. 23, no. 12, Dec. 2005.

[7] A. Zemlianov and G. de Veciana, "Cooperation and decision-making in a wireless multi-provider setting," in *Proc. IEEE INFOCOM*, 2005, pp. 1–14.

[8] S. Shakkottai, E. Altman, and A. Kumar, "The case for non-cooperative multihoming of users to acess points in IEEE 802.11 WLANs," in *Proc. IEEE INFOCOM*, Barcelona, Spain, Apr. 2006.

[9] E. Altman and T. Başar, "Optimal rate control for high speed telecommunication networks," in *Proc. of the 34th IEEE Conference on Decision and Control*, New Orleans, LA, December 1995, pp. 1389–1394.

[10] J. P. Singh, T. Alpcan, P. Agrawal, and V. Sharma, "An optimal flow assignment framework for heterogeneous network access," submitted to IEEE Intl. Symp. on a World of Wireless, Mobile and Multimedia Networks, 2007.

[11] J. Navratil and R. L. Cottrell. Abing. [Online]. Available: http://www-iepm.slac.stanford.edu/tools/abing/

[12] T. Başar and P. Bernhard, *H$^{\infty}$-Optimal Control and Related Minimax Design Problems: A Dynamic Game Approach*, 2nd ed. Boston, MA: Birkhäuser, 1995.