

Game Theory for Heterogeneous Flow Control

Ao Tang* and Lachlan L. H. Andrew†

*School of Electrical and Computer Engineering
Cornell University, Ithaca, NY 14853, USA

†Division of Engineering and Applied Science
California Institute of Technology, Pasadena, CA 91125, USA

Abstract—A general framework is developed for networks with flows that use all available congestion signals to regulate their rates. It is conceptually a generalization of the existing network utility maximization (NUM) theory for homogeneous congestion control. Instead of a convex optimization characterization in NUM, a game with multiple convex optimizations is formulated to characterize equilibria in such a network. Examples are provided to motivate the needs of this general theory. We also provide some basic properties of the game and point out some possible future directions along this line.

I. INTRODUCTION

Various proposals that use different congestion signals besides packet loss have been proposed. However, the standard theory which is based on network utility maximization (NUM), e.g. [7], [10], [9], [11], relies critically on the assumption that all flows use the *same* congestion signal and therefore no longer works in a heterogeneous network. Recently, a systematic study has been carried out to analyze all major properties of equilibrium in a heterogeneous network [20], [21] with the assumption that each flow uses a *single* but could be *different* congestion signal to adapt its rate. Conceivably, a more general setting should allow flows to use all available congestion signals, or any subset thereof, to regulate their rates. This not only provides more information on congestion to sources but also is a better model as some existing protocols [3], [8], [17], [18], including the currently used descendants of standard TCP Reno [5], [12], which actually respond to both packet loss and delay and therefore use multiple congestion indications. This paper is devoted to setting up such a framework and we will show that instead of an optimization model in NUM, a game is needed to characterize the equilibria of such networks. Both the standard network utility maximization and the recent work on heterogeneous congestion control then become special or limiting cases of this general framework.

The paper is organized as follows. We start with a brief review of the existing NUM theory with an emphasis on its predictions (section II). This will be immediately followed by motivating examples which show limitations of these predictions (section III). In section IV, we then define network equilibrium in networks with flows using multiple congestion signals. This is done by including the physical relations among various congestion signals. A general mathematical framework is considered in section V by relaxing those physical constraints. We conclude (section VI) by discussing some possible important extensions.

II. EXISTING THEORY AND ITS PREDICTIONS

When each link has a unique price, and all sources respond to the sum of the prices of links on their paths, then the equilibrium is the unique solution of the following utility maximization problem defined in [7] and its Lagrange dual [10]:

$$\max_{x \geq 0} \sum_i U_i(x_i) \quad (1a)$$

$$\text{subject to } Rx \leq c \quad (1b)$$

where c is the capacity vector, x is the rate vector and R is the routing matrix. the utility functions (U_i 's) are derived from the equilibrium equations of the congestion control protocols, and are typically increasing and strictly concave. The more rigorous definitions of these notations will be provided in section IV. In general, the compactness of the feasible set and the strict concavity of U_i guarantees the existence and uniqueness of the optimal solution of (1a)–(1b).

The basic idea to relate the utility maximization problem (1a)–(1b) to the equilibrium of the flow control algorithm is to examine the dual of the utility maximization problem, and interpret the congestion signal p_l generated by link l as a Lagrange multiplier associated with each link capacity constraint (see, e.g., [9], [10]). As long as the congestion measure p_l increases whenever a link is over-utilized and decreases when $p_l > 0$ and the link is underutilized, then the only equilibrium values of p_l will be the Lagrange multipliers, regardless of the specific dynamics of the AQM mechanism.

The utility maximization problem provides a compact and global characterization of the whole system. In particular, the following two claims are frequently used as examples of its corollaries.

- The equilibrium congestion signals (dual variable) and therefore the equilibrium rates (primal variables) are independent of the AQM settings.
- The system always admits a unique equilibrium.

In [20], [21], it was shown that, when different flows use different types of congestion signals, the utility maximization framework breaks down and the corollaries above do not hold. In this paper, we proceed and allow all sources to use multiple congestion signals simultaneously.

In the next section, we will use standard TCP New Reno protocol and show that even for a network with only Reno protocol, the two above corollaries do not apply.




Fig. 1. Topology used for Example 1. Two flows of different RTTs share a link, with variable buffer size B .

III. MOTIVATING EXAMPLES

The following two examples demonstrate that, even with a single congestion control algorithm, it is possible for the two primary claims above to be false, if the algorithm responds to two types of congestion indication. The examples will use an idealized version of TCP Reno, whose window size,

$$W_i = \sqrt{\frac{2}{p_i}} \quad (2)$$

depends purely on the long-term average packet loss rate it experiences, p_i . (This model is like that of [9], except that it ignores a term only relevant when W_i is small.) Its equilibrium rate is then

$$x_i = \frac{1}{\tau_i + d_i} \sqrt{\frac{2}{p_i}}, \quad (3)$$

where τ_i , is the total queueing delay it experiences, and d_i is the round trip time (RTT) in the absence of queueing. Clearly, x_i depends on two separate congestion measures, p_i and τ_i .

A. Dependence of equilibrium on AQM parameters

In traditional queueing networks, increasing the buffer size reduces the packet loss rate. The addition of utility maximization flow control makes the loss rate less sensitive to AQM parameters, but not entirely independent.

This is easily seen by considering a bottleneck link with capacity C and buffer size B packets, carrying two flows, one with d_1 and one with $d_2 = kd_1$, as shown in Figure 1. With AQM parameter $B \ll Cd_1$, the rates are approximately in the ratio $x_1/x_2 = k/1$, and the packet loss rate is $p \approx 2/(2 + (kd_1C/(k+1))^2)$, regardless of the precise value of B . However, for large changes in the AQM settings this no longer holds. As $B \rightarrow \infty$, the secondary congestion indicator τ (delay) increases causing both the primal and dual variables to change, with $x_1, x_2 \rightarrow C/2$ and $p \rightarrow 0$.

This is illustrated in Figures 2 for the case of $d_1 = 30$ ms, $d_2 = 100$ ms and $C = 8333$ pk/s (100 Mbit/s with 1500 byte packets). The solid lines show the rates (primal variables) achieved by the two flows as the buffer size increases. For small buffers, doubling the buffer size has negligible impact, but when the buffering becomes large enough to affect the RTT, the short-RTT flow reduces its rate in response to the congestion. The packet loss rate (dual variable) is also almost independent of the buffer size when the buffer is small, but depends strongly on in when it becomes larger.




Fig. 2. Example 1. Solid lines: Rates of two flows with different propagation delays sharing a bottleneck, as a function of buffer size. Dashed line: Corresponding packet loss probability.




Fig. 3. Topology used for Example 2. If the two-hop flows have high rate, link 2 is a bottleneck. If the three-hop flows have high rate, links 1 and 3 are bottlenecks.

B. Multiple equilibria

The following construction is analogous to that for heterogeneous congestion control presented in [19].

Consider the network shown in Figure 3, consisting of three links, with capacities $C_1 = C_3 = 8333$ packets/second (100 Mbit/s) and with buffer sizes $B_1 = B_3 = 1000$ packets and $B_2 = 0$ packets. Consider also two-hop TCP flows with $d_1 = d_3 = 3$ ms, traversing links 1 and 2 and links 2 and 3 respectively, and two three-hop TCP flows traversing all links, each with $\tau_2 = 100$ ms. By symmetry, the two three-hop flows have equal rate x_2 , and the two two-hop flows have equal rate x_1 .

For certain C_2 , this network can have two equilibria, one in which there is a bottleneck only at link 2, and one in which there are bottlenecks at links 1 and 3, but not at link 2. This will be explained with the help of Figure 4, which shows the fraction of unused capacity on the non-bottleneck link as C_2 varies. Specifically, the dashed line shows $(C_1 - x_1 - 2x_2)/C_1$ when link 2 is a bottleneck ($2x_1 + 2x_2 = C_2$), and the solid line shows $(C_2 - 2x_1 - 2x_2)/C_2$ when link 1 is a bottleneck ($x_1 + 2x_2 = C_1$).

When $C_2 = C_1$, link 2 must be a bottleneck because it carries more traffic. Thus there is only the dashed line at that point in the graph. In this case, the RTT of the two-hop flows




Fig. 4. Example 2: The equilibrium proportion of unused capacity on the non-bottleneck link(s) as the ratio of capacities varies, for $C_1 = 8333$ pk/s, $d_1 = 3$ ms, $d_2 = 100$ ms and $B = 1000$ packets. For ratios between about 1.45 and 1.9, two equilibria exist: one with link 2 as a bottleneck and one with links 1 and 3 as bottlenecks.

is only d_1 since there is no buffer at link 2. As C_2 increases, the rates all increase reducing the spare capacity on link 1, but maintaining the small RTTs. Similarly, when $C_2 = 2C_1$, links 1 and 3 must be bottlenecks, since they carry over half as much traffic as link 2. In this case, the bottleneck buffers are full, and the two-hop flows have RTT significantly above d_1 .

For C_2/C_1 between about 1.45 and 1.9, both situations are possible. If links 1 and 3 are initially bottlenecked, the large queueing delay will cause x_1 to remain small and so they will remain bottlenecked. The queueing delay has less effect on x_2 because $d_2 \gg d_1$. Conversely, if link 2 is initially bottlenecked, the high loss rate will cause x_2 to remain small and link 2 will remain bottlenecked.

These two examples clearly show that even in the case when there is only one type of protocol (TCP Reno in these examples), the standard utility maximization framework still may not apply if the protocol uses multiple congestion signals. In the following sections, we provide an alternative framework to understand such networks.

IV. NETWORK EQUILIBRIUM

A network consists of a set L of links indexed by $l \in L$, a set I of flows indexed by $i \in I$ and a set J of types of congestion signals, indexed by $j \in J$. Let R be the $L \times I$ routing matrix: $R_{li} = 1$ if source i uses link l and 0 otherwise. Each link l has congestion signals p_l^j , and for each $j \in J$, each flow is notified of the sum of the type j congestion signals of flows on its path.

Each link emits multiple congestion signals, such as packet loss and queueing delay, to which algorithms may respond. However, the congestion signals produced by a single link are closely related. As in [20], we model this relationship through a price mapping function that maps a common price (e.g., queue length) at a link to different prices (e.g., loss probability

and queueing delay) observed by different sources. Formally, every link l has a price p_l . The type j congestion signal is the sum of the “effective prices” $m_l^j(p_l)$ on a path, where m_l^j is a price mapping function, which can depend on both the link and the congestion signal type. The exact form of m_l^j depends on the AQM algorithm used at the link; see [19] for links with RED.¹ Let $m^j(p) = (m_l^j(p_l), l = 1, \dots, L)$. The aggregate type j price for source i is defined as

$$q_i^j = \sum_l R_{li} m_l^j(p_l) \quad (4)$$

Let $q^j = (q_i^j, i = 1, \dots, N^j)$. Then $q^j = R^T m^j(p)$.

Let x be a vector of source rates, x_i . In general, if z_k are defined, then z denotes the (column) vector $z = (z_k, \forall k)$. Other notations will be introduced later when they are encountered.

Each flow $i \in I$ sets its rate x_i in response to all of the congestion signals. At equilibrium,

$$x_i = f_i^0(q_i^1, q_i^2, \dots, q_i^J) \quad (5)$$

Here, f_i^0 only depends on the TCP algorithm of flow i .

As usual, we use $x(q) = x(q^1, q^2, \dots, q^J)$ to denote the vector-valued functions composed of x_i . Since $q = R^T m(p)$, we often abuse notation and write $x_i(p)$ and $x(p)$.

Define the aggregate source rates $y(p) = (y_l(p), l = 1, \dots, L)$ at links l as:

$$y(p) = Rx(p) \quad (6)$$

In equilibrium, the aggregate rate at each link is no more than the link capacity, and the link price can be strictly positive only when the aggregate rate equals the capacity at that link. Formally, we call p an *equilibrium price*, a *network equilibrium*, or just an *equilibrium* if it satisfies (from (4)–(6))

$$P(y(p) - c) = 0, \quad y(p) \leq c, \quad p \geq 0 \quad (7)$$

where $P := \text{diag}(p_l)$ is a diagonal matrix. The network equilibrium is specified by (4)–(7). Let E be the network equilibrium set, we have:

$$E = \{p \in \mathfrak{R}_+^L \mid P(y(p) - c) = 0, \quad y(p) \leq c\} \quad (8)$$

Throughout this paper, we also adopt some standard assumptions.

A1: Price mapping functions m_l^j are continuously differentiable in their domains and non-decreasing with $m_l^j(0) = 0$ for all j, l . For each p_l , at least one m_l^j is strictly increasing.

A2: The demand function f_i^0 is differentiable and strictly decreasing in each variable. Moreover, assume that for any q^{-j} with $q_i^k > 0$ for all k , $\lim_{q_i^j \rightarrow \infty} f_i^0 = 0$.

These are mild assumptions. The first assumption on m_l^j preserves the relative order of prices and maps zero price to zero effective price. Assumption A2 says that, if one kind of price is large enough, provided others are not zero, then the rate can be constrained to be arbitrarily small.

¹One can also take one particular type of price p_l^j , e.g., queueing delay, as the common price p_l . In this case the corresponding price mapping function is the identity function, $m_l^j(p_l) = p_l$.

V. A GAME THEORETICAL CHARACTERIZATION

In this section, we will show that the network equilibrium which was defined in last section is a subset of the set of Nash equilibrium of a J -persons convex game. The basic idea is the same as in the homogeneous case where we ignore the physical constraints that are imposed by link AQMs and focus on the mathematical optimization problem that is determined by the equilibrium of end to end TCPs.

Let $q_i^{-j} = (q_i^1, \dots, q_i^{j-1}, q_i^{j+1}, \dots, q_i^J)$. By the implicit function theorem, (5) may be rearranged as

$$q_i^j = f_i^{-j}(x_i, q_i^{-j}), \quad (9)$$

where f_i^{-j} is decreasing in x_i .

Define partial utility functions

$$U_i^j(x_i, q_i^{-j}) = \int f_i^{-j}(x_i, q_i^{-j}) dx. \quad (10)$$

Note that U_i^j is strictly concave increasing, because f_i^{-j} is decreasing.

Take the standard TCP Reno as an example. Recall from (3) that its equilibrium rate can be modeled as

$$x_i = f_i^0(p_i, \tau_i) = \frac{1}{\tau_i + d_i} \sqrt{\frac{2}{p_i}} \quad (11)$$

where d_i is the fixed ‘‘propagation’’ delay and the congestion signals are τ_i , the total queueing delay experienced, and p_i , the packet loss probability. Identifying $(q^1, q^2) \equiv (p_i, \tau_i)$, its equilibrium satisfies

$$q_i^1 = p_i = f_i^{-1}(x_i, \tau_i) = \frac{2}{x_i^2(\tau_i + d_i)^2} \quad (12)$$

$$U_i^1(x_i, \tau_i) = -\frac{2}{x_i(\tau_i + d_i)^2} \quad (13)$$

$$q_i^2 = \tau_i = f_i^{-2}(x_i, p_i) = \frac{1}{x_i} \sqrt{\frac{2}{p_i}} - d_i \quad (14)$$

$$U_i^2(x_i, p_i) = \sqrt{\frac{2}{p_i}} \log(x_i) - x_i d_i. \quad (15)$$

We now define the following game. There are J players, with the j th player able to choose p_l^j for each link l , subject to the feasibility constraint

$$Rf_i^0(q_i) \leq c. \quad (16a)$$

The payoff for the j th player is the sum of the type j utility functions for all flows under p^j :

$$\sum_i U_i^j(f_i^0(q_i), q_i^{-j}) \quad (16b)$$

In other words, for each $j \in \{1, 2, \dots, J\}$, the j th player tries to solve

$$\max_{x \geq 0} \sum_i U_i^j(x_i; q_i^{-j}) \text{ subject to } Rx \leq c \quad (17)$$

Given that this is a convex optimization with Slater’s condition [1] satisfied, strong duality holds and we can also equivalently look at its dual [10]:

$$\min_{p^j \geq 0} \sum_i \max_{x_i \geq 0} \left(U_i^j(x_i; q_i^{-j}) - x_i \sum_l R_{li} p_l^j \right) + \sum_l c_l p_l^j. \quad (18)$$

Equation (17), or equivalently (18), defines the noncooperative game we are studying in this paper. The remainder of this section is devoted to setting up some basic results of this game.

Theorem 1. *All network equilibria are Nash equilibria of the corresponding game defined by (17) or (18).*

Proof: By the construction of U_i^j and the optimality condition of (17), equations (5) to (7) define the Nash equilibria of the game. Thus any network equilibrium, which by definition satisfies (4)–(7), is a Nash equilibrium of the game. ■

Theorem 2. *The game always admits at least one pure Nash equilibrium.*

Proof: The joint strategy set² is nonempty, convex and compact subset of a Euclidian space. The utility of each player is continuous and concave on its action set. By the theorem of Rosen [14], the game has a pure Nash equilibrium. ■

Remark: In the standard flow control where there is only one type of congestion signal, a convex optimization problem characterizes the equilibrium. It provides a unique equilibrium and the corresponding congestion signals. If the congestion signal p_l is within the dynamic range which can be physically generated by the network AQM, then it is also the network equilibrium.³ This is illustrated in Figure 5.

Now, we have J convex optimization problems that are tangled together and their Nash equilibrium set includes all the possible network equilibria. Again like the single congestion measure case, we need to check the AQM conditions (4) to find the final network equilibria. This relationship is illustrated in Figure 6.

VI. CONCLUSION

This paper sets up a general framework for network congestion control where each flow decides its rate based on possibly more than one congestion signal. It is shown that unlike in the single congestion signal case where a convex optimization fully characterizes the possible equilibrium, a game involves multiple convex optimizations is needed and the saddle points of the game include the set of network equilibria.

The study is still preliminary and many problems remain open. It is of particular interest to find equilibria of the game

²Here, individual strategy set may depend on strategies of other players. Therefore the standard Nash existence theorem [13] does not apply.

³This need not be the case; algorithms such as Vegas [2] and FAST [22] which respond to delay have an upper bound placed on p_l by the size of their buffers, while algorithms such as Reno which respond to loss may have a lower bound placed on p_l by physical-layer packet corruption, or by infrequent but intense bursts of cross traffic.




Fig. 5. The optimization formulation (1) has an equilibrium regardless of the AQM. Only a subset of AQM settings actually support the equilibrium in a network




Fig. 6. The game formulation (16) has an equilibrium regardless of the AQM. Only a subset of equilibria actually have prices which can be implemented by the AQM, satisfying $p_l^j = m^j(p_l)$ for all l and j for some p_l .

described by (17) or (18) and of the network itself. This requires investigation about the global dynamical behaviors of such systems. It is challenging given the possibility of having multiple equilibria. We expect such studies, if fruitful, will interest networking community as well as several other communities such as dynamical systems, microeconomics and theoretical computer science. Existing closely related work such as monotone dynamical systems [16], global Newton method [15], S-modular games [6] and general Nash equilibrium computation [4] may help find the right tools. Special structures of the problem, e.g., the relations among different U_i^j 's which are all derived from the same equilibrium condition, are likely to be critical for obtaining stronger than usual general results.

REFERENCES

- [1] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [2] L. S. Brakmo and L. L. Peterson. TCP Vegas: End-to-End Congestion Avoidance on a Global Internet, *IEEE Journal on Selected Areas in Communication*, Vol. 13, Nov. 8, 1995.
- [3] C. Casetti, M. Gerla, S. Mascolo, M. Sansadidi and R. Wang. TCP Westwood: end to end congestion control for wired/wireless networks. *Wireless Networks Journal*, vol.8, pp.467-479, 2002.
- [4] A. Fabrikant, C. Papadimitriou and K. Talwar. The Complexity of Pure Nash Equilibria. *Proceedings of STOC*, 2004.
- [5] S. Floyd and T. Henderson. The NewReno Modification to TCP's Fast Recovery Algorithm. RFC 2582, 1999.
- [6] D. Fudenberg and J. Tirole *Game Theory*. The MIT Press, 1991.
- [7] F. Kelly, A. Maoullou, and D. Tan. Rate control for communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research Society*, 49:237-252,1998.
- [8] S. Liu, T. Basar and R. Srikant. TCP-Illinois: A loss and delay-based congestion control algorithm for high-speed networks. *Proc. First Int. Conf. on Perform. Eval. Methodol. Tools (VALUETOOLS)*, 2006.

- [9] S. Low. A duality model of TCP and queue management algorithms. *IEEE/ACM Transactions on Networking*, 11(4):525-536, August 2003.
- [10] S. Low and D. Lapsley. Optimization flow control, I: basic algorithm and convergence. *IEEE/ACM Transactions on Networking*, 7(6): 861-874, December 1999.
- [11] S. Kunniyur and R. Srikant. End-to-end congestion control: utility functions, random losses and ECN marks. *IEEE/ACM Transactions on Networking*, 11(5):689-702, Oct 2003.
- [12] M. Mathis and J. Mahdavi and S. Floyd and A. Romanow. TCP Selective Acknowledgement Options. RFC 2018, 1996.
- [13] M. Osborne and A. Rubinstein. *A Course in Game Theory*. The MIT Press, 1994.
- [14] J. Rosen. Existence and Uniqueness of Equilibrium Points for Concave N-Person Games *Econometrica*, 33(3):520-534, 1965.
- [15] S. Smale. A convergent process of price adjustment and global newton methods. *Journal of Mathematical Economics*, 3(2):107-120, July 1976.
- [16] H. Smith. *Monotone Dynamical Systems: An Introduction to the Theory of Competitive and Cooperative Systems*. American Mathematical Society, Mathematical Surveys and Monographs, 1995.
- [17] K. Tan, J. Song, Q. Zhang, and M. Sridharan. A compound tcp approach for high-speed and long distance networks. *Proceedings of IEEE Infocom*, April 2006.
- [18] A. Tang, D. Wei, S. Low and M. Chiang. Heterogeneous Congestion Control: Efficiency, Fairness and Design. *Proceedings of IEEE ICNP*, 2006.
- [19] A. Tang J. Wang S. Hegde S. H. Low. Equilibrium and Fairness of Networks Shared by TCP Reno and Vegas/FAST *Telecommunications Systems*, 30(4):417-439, Dec 2005.
- [20] A. Tang, J. Wang, S. Low and M. Chiang. Equilibrium of Heterogeneous Congestion Control: Existence and Uniqueness. *IEEE/ACM Transactions on Networking*, 15(4):824-837, August 2007.
- [21] A. Tang, X. Wei, S. Low and M. Chiang. Equilibrium of Heterogeneous Congestion Control: Optimality and Stability. Submitted to *IEEE/ACM Transactions on Networking*, September 2007.
- [22] D. Wei, C. Jin, S. H. Low, and S. Hegde. FAST TCP: motivation, architecture, algorithms, performance. *IEEE/ACM Transactions on Networking*, 14(6):1246-1259, Dec. 2006.