

Can TCP Metastability Explain Cascading Failures and Justify Flow Admission Control in the Internet?

V. Marbukh

National Institute of Standards and Technology
100 Bureau Drive
Gaithersburg, MD 20899 USA
E-mail: marbukh@nist.gov

Abstract—This paper discusses implications of possible metastability of TCP-type fair bandwidth sharing under random flow arrivals/departures for understanding and defending the Internet against cascading failures. Cascading failures can be viewed as a process of network transition from desirable metastable mode with finite number of flows in progress to the congested mode. It is possible to eliminate or reduce the possibility of cascading failures with properly designed flow admission control which stabilizes the network in a close neighborhood of the desirable metastable state.

Keywords—Internet, TCP, metastability, cascading failures, flow admission control.

I. INTRODUCTION

Flow-level abstraction for fair bandwidth sharing among fixed set of flows has been proposed in [1] and a Markov model for flow-level performance of bandwidth sharing protocols under random flow arrivals/departures has been proposed in [2]-[3]. The Markov model assumes separation of time scales: the bandwidth sharing protocol reaches equilibrium bandwidth sharing much faster than the numbers of flows in progress change due to flow arrivals/departures. However, the high dimension of the corresponding Markov process makes this Markov model intractable for practical-size networks. The difficulty results from a large number of interacting elements (processes) describing numbers of flows with different source-destinations, where the interactions occur due to sharing link bandwidth by different flows. In the fluid asymptotic regime [4] the numbers of flows in progress becomes deterministic, making fluid model tractable.

It has been observed by simulation in [5] that packet retransmissions during file transfers may cause instability of fair bandwidth sharing under arriving/departing flows when Explicit Congestion Notification (ECN) signals are represented by packet losses. A goodput-based Markov model, which accounts for packet retransmissions by assuming that file transmission rates are determined by the end-to-end goodputs rather than the corresponding throughputs, has been proposed in [6]. The goodput-based model is typically unstable for a network with multihop

routing even under light load when the corresponding throughput-based model is stable. This instability is the result of the positive feedback created by increase in the link packet losses with increase in the number of flows carried on the link. Despite this instability, a desirable metastable, i.e., persistent, network state with finite numbers of flows in progress may still exist for sufficiently light exogenous load [6]. An appropriately designed flow admission control can transform the desirable metastable state into the unique stable state at the cost of small flow rejection probability.

This paper suggests that these theoretical results may contribute to understanding of the cascading failures in the Internet and how to defend the Internet against such failures. We consider TCP Reno type fair bandwidth sharing, when ECN signals are delivered to the sources through end-to-end packet losses. One may assume that due to significant Internet underutilization, metastability rather than stability of the observed Internet state may not be apparent. Utilization increase due to physical infrastructure failures or load increase as a result of a fluctuation, denial of service attacks, or flash crowds may “push” the network out of the desirable metastable state initiating process of cascading failures and leading to a congested mode.

The paper is organized as follows. Section 2 briefly introduces a fixed-point performance model of a fair bandwidth sharing with fixed set of flows [7]. Section 3 describes Markov [2]-[3] and fluid [4] performance models of fair bandwidth sharing [1] under arriving/departing flows. Section 4 discusses stability/instability of these Markov and fluid models and the importance of flow admission control. Section 4 also suggests possible implications of these results for understanding cascading failures in the Internet and defending the Internet against cascading failures. Section 5 considers specific case of a symmetric ring network. Finally, Conclusion summarizes our results and outlines directions for future research.

II. FIXED-POINT MODEL FOR A FIXED SET OF FLOWS

Consider a network with links $j \in J$. Given a directed route $r = \{j_1, \dots, j_m, j, j_{m+2}, \dots, j_M\}$ passing through link

j let $r_j^- = (j_1, \dots, j_m)$ be the part of the route r lying downstream from link j , and $r_j^+ = (j_{m+2}, \dots, j_M)$ be the part of route r lying upstream from link j . Given a vector of numbers of flows on all feasible routes $r \in R$: $n = (n_r)$, this section briefly describes models for evaluating the vector of end-to-end throughputs $x = (x_r)$ and good-puts $g = (g_r)$.

Weighted (α, w) -fair bandwidth allocation [1] assigns the transmission rate on a route $r \in R$ as follows:

$$x_r = \left(w_r / \sum_{j \in r} p_j \right)^{1/\alpha}, \quad (1)$$

where fixed parameters $\alpha, w_r > 0$ characterize fairness, and parameters p_j represent Explicit Congestion Notification (ECN) on links $j \in J$. Further, in this paper we assume that p_j is the packet loss probability on link $j \in J$. When $w_r = 1$, the cases $\alpha \rightarrow 0$, $\alpha \rightarrow 1$ and $\alpha \rightarrow \infty$ correspond respectively to an allocation which achieves maximum throughput, is proportionally fair, or is max-min fair. Case $\alpha = 2$ corresponds to TCP-Reno rate assignment (1) in a case of small packet losses

$$p_r \approx \sum_{j \in r} p_j < 1. \quad (2)$$

Note that a more general rate allocation

$$x_r = (w_r / p_r)^{1/\alpha} (1 - p_r)^\beta \quad (3)$$

with some $\beta \geq 0$ can be used to describe a case when $x_r \downarrow 0$ as $p_r \uparrow 1$. The case $\alpha = \beta = 2$ has been considered in [8].

Under the assumption of independent packet losses on different links, the end-to-end packet loss probability p_r is

$$p_r = 1 - \prod_{j \in r} (1 - p_{rj}) \quad (4)$$

and the load on a link j resulted from a flow carried on a route r is

$$x_{rj} = x_r \prod_{i \in r_j^-} (1 - p_{ri}), \quad (5)$$

where p_{ri} is the probability that a packet from a flow carried on a route r is lost on a link i , and x_r is the source transmission rate on route r .

Given a mixture of flows in progress on different routes $N = (N_r)$, we assume that the packet loss probability on a link j for a flow carried on a route r is

some known function of the average bandwidth allocated to each flow on the link j :

$$p_{rj} = p_{rj}(N_r, x_{rj}, \forall r' : j \in r' \subset R) \quad (6)$$

Assuming that functions (3) and (6) are known, system (3)-(6) forms a closed system of fixed-point equations for the end-to-end performance (x, p) , given a vector of numbers of flows in progress on different routes $N = (N_r, \forall r \in R)$. Here the vector of source transmission rates is $x = (x_r, \forall r \in R)$, and the vector of link packet losses for different sources is $p = (p_{rj}, \forall j \in J)$.

Note that while functional form (3) depends on the specific TCP implementation, functional form (6) depends on the buffer queue management discipline. In this paper we assume that link packet loss probability is the same for all flows carried on the link:

$$p_{rj} = p_j, \quad \forall r : j \in r \subset R \quad (7)$$

Assumption (7) holds in a bufferless case as well as in a case of buffered links with First In First Out (FIFO) packet scheduling. In a case (7)

$$p_j \geq \tilde{p}_j \stackrel{def}{=} \max \left\{ 0, 1 - 1 / \sum_r N_r x_{jr} \right\} \quad (8)$$

where link packet loss probability $\tilde{p}_j$ corresponds to the fluid limit when link capacity and the number of flows carried on the link are large.

Given a mixture of flows in progress $N = (N_r)$, the Fixed Point Model determines the end-to-end packet loss probability on a route r

$$p_r = 1 - \prod_{j \in r} (1 - p_j), \quad (9)$$

and the corresponding good-put

$$g_r = x_r \prod_{j \in r} (1 - p_j). \quad (10)$$

The aggregate good-put on a route r is

$$G_r = N_r g_r = N_r x_r \prod_{j \in r} (1 - p_j) \quad (11)$$

III. MODELS FOR ARRIVING/DEPARTING FLOWS

Following [4] assume that the network carries two types of flows: file transfers and streaming flows. Introduce vector (N_1, N_2) , where $N_1 = (N_{1r})$ and $N_2 = (N_{2r})$ are the vectors of the numbers of file transfers and streaming flows, respectively, carried on all feasible routes $r \in R$. We assume that files arrive on a route $r \in R$ according to a Poisson process of rate $\Lambda_{1r} = \Lambda_{1r}(N_1, N_2)$, and streaming flows arrive on route r according to a Poisson process of rate

$\Lambda_{2r} = \Lambda_{2r}(N_1, N_2)$, where dependence on the vector of flows in progress (N_1, N_2) allows for modeling state-dependent admission control and routing as well as user dissatisfaction with deteriorating good-put as the network becomes more congested. The size of a file transmitted on a route $r \in R$ is distributed exponentially with average b_r . The holding time of a streaming flow carried on a route $r \in R$ is distributed exponentially with average θ_r . All arrival processes, file sizes and durations of streaming flows are jointly statistically independent.

We assume that flow control operates on much faster time-scale than the process of flow arrivals/departures. Given a vector of flows in progress (N_1, N_2) , this assumption of separation of time-scales allows us to approximate the conditional link losses by the solution to the corresponding fixed point system (3)-(6) $p(N) = (p_j(N))$ with a fixed number of flows in progress $N = N_1 + N_2$. Under our assumptions, the vector $(N_1(t), N_2(t))$ is a homogeneous in time $t \in [0, \infty)$ vector Markov process. Probabilities $P(t, N_1, N_2) = \Pr\{N_1(t) = N_1, N_2(t) = N_2\}$ are uniquely determined by the following system of linear differential Kolmogorov equations:

$$\begin{aligned} & \dot{P}(t, N_1, N_2) = \\ & -P(t, N_1, N_2) \sum_r [\Lambda_{1r}(N) + \Lambda_{2r}(N)] \\ & -P(t, N_1, N_2) \sum_r [b_r^{-1} g_r(N) N_{1r} + \theta_r^{-1} N_{2r}] \\ & + \sum_r P(t, N_1 - 1_r, N_2) \Lambda_{1r}(N - 1_r) \\ & + \sum_r P(t, N_1, N_2 - 1_r) \Lambda_{2r}(N - 1_r) \\ & + \sum_r P(t, N_1 + 1_r, N_2) (N_{1r} + 1) g_r(N + 1_r) b_r^{-1} \\ & + \sum_r P(t, N_1, N_2 + 1_r) (N_{2r} + 1) \theta_r^{-1} \end{aligned} \quad (12)$$

$$N = N_1 + N_2 \quad (13)$$

supplemented with the normalization condition

$$\sum_{(N_1, N_2)} P(t, N_1, N_2) \equiv 1, \quad (14)$$

given initial conditions $P(0, N_1, N_2)$. In (6), vector 1_r has all components equal zero, except for the component corresponding to route r , which is equal to 1.

If the process $(N_1(t), N_2(t))$ is ergodic, i.e., unique proper steady state distribution

$$P(N_1, N_2) = \lim_{t \rightarrow \infty} P(t, N_1, N_2) \quad (15)$$

exists, it satisfies the corresponding system of linear algebraic Kolmogorov steady-state equations, which describe the equilibrium of non-steady Kolmogorov equations (12). Steady-state performance criteria can be expressed in terms of distribution (15), e.g., according to the Little formula the average file transfer time on route $r \in R$ in a case when these files arrive at a constant rate Λ_{1r} is

$$\tau_r = \frac{1}{\Lambda_{1r}} \sum_{N_{1r}=0}^{\infty} N_{1r} P(N_1, N_2) \quad (16)$$

However, due to difficulty of solving steady-state Kolmogorov equations for practical-size networks, this direct approach to network performance evaluation is computationally infeasible. The rest of this section briefly introduces the fluid approximation.

Consider the fluid asymptotic regime, when link capacities C_j and flow arrival rates Λ_{ir} are scaled up proportionally to a large parameter ε^{-1} , while the average file sizes b_r and average durations of streaming flows θ_r are kept constant:

$$\begin{aligned} C_j &= c_j \varepsilon^{-1}, \quad \Lambda_{ir} = \lambda_{ir} \varepsilon^{-1}, \quad \varepsilon \rightarrow 0, \\ c_j, \lambda_{ir} &= O(1), \quad i=1,2; \forall j \in J, \forall r \in R \end{aligned} \quad (17)$$

It is possible to show [4] that under the asymptotic regime (17) typical numbers of flows carried on all feasible routes $r \in R$, N_{ir} , are also scaled up proportionally to the large parameter ε^{-1} :

$$N_{ir} = n_{ir} \varepsilon^{-1}, \quad n_{ir} = O(1), \quad \varepsilon \rightarrow 0, \quad i=1,2; r \in R \quad (18)$$

where vector $n_i = (n_{ir}(t), r \in R)$ evolution on any finite time interval $t \in (0, T]$, $0 < T < \infty$ is described by the following system of ordinary differential equations:

$$\dot{n}_{1r} = \lambda_{1r} (n_1 + n_2) - b_r^{-1} n_{1r} g_r(n_1 + n_2) \quad (19)$$

$$\dot{n}_{2r} = \lambda_{2r} (n_1 + n_2) - \theta_r^{-1} n_{2r} \quad (20)$$

IV. METASTABILITY AND CASCADING FAILURES

Informally, bandwidth allocation is stable under arriving/departing flows if the number of flows in progress remains finite as time progresses. However, even if the numbers of flows in progress infinitely grows with time $t \rightarrow \infty$, the number of flows in progress may not grow on a time scale of practical interest. This possibility arises for large-scale networks described by a multi-component Markov process $(N_1(t), N_2(t))$ with a large number of components. Generally speaking convergence time to the steady-state distribution for an ergodic Markov process with finite number of states depends on the number of

states. As the number of states increases, the convergence time may infinitely increase creating the possibility of metastable, i.e., persistent states. If the “life span” of the corresponding metastable state is longer than the time scale of practical interest, this metastable state from a practical perspective appears to be stable.

In a case of flow arrival rates independent of the number of flows in progress: $\Lambda_{ir}(N) \equiv \Lambda_{ir}$, the necessary and sufficient conditions for ergodicity of the conventional “throughput-based” Markov model with file transfer rates determined by the end-to-end throughputs

$$\rho_{1j} \stackrel{\text{def}}{=} \frac{1}{C_j} \sum_{r: j \in r} \Lambda_{ir} b_r < 1, \quad \forall j \in J \quad (21)$$

have been obtained in [2]-[4]. Conditions (21) have a simple interpretation: aggregate link loads are less than the link capacities.

However, the “goodput-based” model is typically unstable even under light load, when conditions (21) are satisfied and the throughput-based model is stable. The instability of the goodput-based model is a result of aggregate goodput (11) deterioration due to increase in the number of packet retransmissions as the number of flows carried on the route links increases. The deterioration is a result of capacity waste: each link $j \in J$ wastes certain portion of its capacity on transmission of “dead” packets to be dropped downstream and then retransmitted by the source. These arguments lead to the following necessary condition for the network stability:

$$\frac{1}{C_j} \sum_{r: j \in r} \frac{\Lambda_{ir}(N) b_r}{\prod_{i \in r_j^+} [1 - p_i(N)]} < 1, \quad \forall j \in J \quad (22)$$

Typically, as the number of flows carried on a link j increases the link packet loss p_j increases approaching unity: $p_j \uparrow 1$, and thus the necessary stability condition (22) is not satisfied for any (arbitrary small) non-zero load in presence of multihop routes. Typical aggregate end-to-end throughput and goodput for file transfer flows are shown in Figure 1.




Figure 1. Goodput deterioration.

Despite instability, a desirable metastable network state with finite number of flows in progress may still exist. One may expect that stability of the network as time progresses $t \rightarrow \infty$ depends on the attractors and topological structure of the system (19)-(20) phase space. It can be shown that in a case when system (19)-(20) has a unique globally asymptotically stable equilibrium point (n_1^*, n_2^*) with finite components the corresponding Markov model is ergodic. Generally, the non-linear system of ordinary differential equations (19)-(20) may have multiple attractors. It is natural to interpret multiple attractors as describing metastable, i.e., persistent, network states. According to this interpretation, changes in the topological structure of the system (19)-(20) describe phase transitions, when network performance characteristics may experience abrupt changes with infinitesimally small change in the exogenous parameters, e.g., load.

In particular, equilibrium points to (19)-(20) are given by the solution to the following system of fixed point equations:

$$n_{1r} = \lambda_{1r}(n_1 + n_2) b_r / g_r(n_1 + n_2) \quad (23)$$

$$n_{2r} = \lambda_{2r}(n_1 + n_2) \theta_r \quad (24)$$

Summarizing equations (23) and (24) we obtain the following fixed point equations for the total average number of flows $n_r = n_{1r} + n_{2r}$ carried on feasible routes $r \in R$:

$$n_r = \lambda_{1r}(n) b_r / g_r(n) + \lambda_{2r}(n) \theta_r \quad (25)$$

In the rest of this section we consider case of flow arrival rates independent of the number of flows in progress: $\lambda_{ir}(n) = \lambda_{ir}$, $i = 1, 2; r \in R$.

We conjecture that system (19)-(20) has the following topological structure. For a sufficiently large initial number of carried flows the number of flows in progress grows infinitely with time: $\|n_r(t)\| \rightarrow \infty$ as $t \rightarrow \infty$. This instability is the manifestation of the non-ergodicity of the Markov model (12)-(14). In absence of streaming flows: $\lambda_{2r} \equiv 0$, system (19) has a locally stable, finite equilibrium $n^* = (n_r^*)$ if and only if

$$\rho_{1j} \stackrel{\text{def}}{=} \frac{1}{C_j} \sum_{r: j \in r} \lambda_{1r} b_r < 1, \quad \forall j \in J \quad (26)$$

Generally, in the presence of streaming flows: $\lambda_{2r} \geq 0$, system (19)-(20) has locally stable finite equilibrium $n^* = (n_r^*)$ if and only if

$$\rho_{1j} < \varphi_j(\lambda_{2r}), \quad \forall j \in J \quad (27)$$

where $\varphi_j(0) = 1$ and the functions $\varphi_j(\lambda_{2r})$ decrease to zero as the streaming flow load λ_{2r} increases. This

equilibrium $n^* = (n_r^*)$ is the only finite attractor of system (19)-(20). The stability domain $\mathbf{N}$ of the equilibrium n^* shrinks with increase in the exogenous load, completely disappearing on the stability boundary of this finite equilibrium.

This behavior of the fluid model suggests that equilibrium n^* corresponds to a metastable, i.e., persistent, state of the Markov model since for sufficiently small exogenous load and number of flows in progress, the end-to-end goodput is sufficient to keep up with the exogenous load. However, eventually the number of flows in progress grows infinitely, since for the sufficiently large number of flows in progress almost the entire network bandwidth becomes occupied with transmitting dead packets to be dropped downstream. A flow admission strategy, which admits flows if and only if the number of flows in progress lies within the stability domain of the equilibrium n^* , transforms the desirable metastable network state into the only stable state at the cost of a small flow blocking probability.

It is known in physics that system transition from a metastable to stable mode occurs in two stages. First, a supercritical embryo of the stable mode appears due to sufficiently large fluctuation of the number of flows in progress or some external event such as physical infrastructure failure, load fluctuation, flash crowds or denial of service attack. Then, this embryo grows following the internal system dynamics, eventually subsuming the entire system. In our case of transition from desirable metastable state to overloaded stable state, this second stage is essentially a process of cascading failures. A flow admission strategy which admits flows if and only if the number of flows in progress lies within the stability domain of the equilibrium n^* modifies the network dynamics eliminating the possibility of cascading failures.

V. EXAMPLE: RING NETWORK

This section briefly considers a case of weighted (α, w) -fair rate allocation (3) with $\beta = 0$ for a K -node ring network, where nodes $k = 1, 2, \dots, K$ and $(k+1) \bmod(K)$ are connected by a directed link j_k of capacity C . We consider a homogeneous case, when weights w_r as well as arrival rates of file transfer flows Λ_{1r} and streaming flows Λ_{2r} are route $r \in R$ independent: $w_r = w$, $r \in R$. Here $\Lambda_{ir} = \Lambda_i$ if $r \in R$ and $\Lambda_{ir} = 0$ otherwise, and set of l -link feasible routes is

$$R = \{(j_k, j_{(k+1) \bmod(K)}, \dots, j_{(k+l) \bmod(K)}) : k = 1, \dots, K\}.$$

We also assume that all feasible routes carry the same number N_1 of file transfer flows as well as the same number N_2 of streaming flows:

$$N_{ir} = \begin{cases} N_i & \text{if } r \in R \\ 0 & \text{otherwise} \end{cases}$$

For a weighted (α, w) -fair rate allocation (3) with $\beta = 0$, source transmission rate is

$$x = \left[\frac{w}{1 - (1-p)^l} \right]^{1/\alpha} \quad (28)$$

Since the load coming on a link j after already traversing k links is thinned by a factor of $(1-p)^k$, we obtain the following expression for the link utilization:

$$\rho = \frac{Nx}{C} \sum_{k=0}^{l-1} (1-p)^k = \frac{Nx}{C} \frac{1 - (1-p)^l}{p}, \quad (29)$$

where the number of flows in progress originated from a given node is N .

Combining (28) with (29) we obtain the following relation between the average link utilization and link packet loss:

$$\rho = \eta \frac{[1 - (1-p)^l]^{1-1/\alpha}}{lp}, \quad (30)$$

where the un-dimensional parameter $\eta = w^{1/\alpha} l N / C$. We consider the fluid regime (17)-(18), when

$$p = \max(0, 1 - 1/\rho), \quad (31)$$

where “effective link utilization” ρ is given by (19). Combining equation (30) and (31), we obtain the following fixed point equation for link packet loss:

$$p = \psi(p) \stackrel{\text{def}}{=} 1 - \frac{lp}{\eta[1 - (1-p)^l]^{1-1/\alpha}} \quad (32)$$

Function $\psi(p)$ is monotonically decreasing for $p \in [0, 1]$, and $\psi(0) = 1$, $\psi(1) < 1$. Equation (32) has unique solution $p^* \in (0, 1)$. It is easy to show the following properties of the solution of equation (32):

$$p^* = l^{-1} \eta^\alpha + o(\eta^\alpha) \text{ as } \eta \rightarrow 0 \quad (33)$$

$$p^* = 1 - l\eta^{-1} + o(\eta^{-1}) \text{ as } \eta \rightarrow \infty \quad (34)$$

In some particular cases equation (32) can be solved explicitly. For example, in a case of proportional fairness, i.e., $\alpha = 1$, equation (32) yields $p^* = (1 + l/\eta)^{-1}$.

After solving equation (32) one can calculate the end-to-end per flow good-put $g = (1-p)^l x$ and the effectiveness of the bandwidth sharing protocol

$$\gamma = Nl g/C. \quad (35)$$

Substituting (29) into (31) we obtain

$$p = 1 - \frac{C}{Nx} \frac{p}{1 - (1-p)^l}. \quad (36)$$

Combining (35) with (36) we obtain the following expression for the effectiveness:

$$\gamma = \frac{lp(1-p)^{l-1}}{1 - (1-p)^l}. \quad (37)$$

Immediately note that in case $l = 1$ formula (37) produces a plausible result $\gamma = 1$, since no bandwidth is wasted on transmission of packets blocked downstream. It is directly follows from (37) that

$$\gamma = 1 - \frac{l-1}{2} p + o(p) \text{ as } p \rightarrow 0 \quad (38)$$

$$\gamma = (1-p)^{l-1} l + o[(1-p)^{l-1}] \text{ as } p \rightarrow 1 \quad (39)$$

Combining (38)-(39) with (33)-(34) we obtain:

$$\gamma = 1 - \frac{l-1}{l} \eta^\alpha + o(\eta^\alpha) \text{ as } \eta \rightarrow 0 \quad (40)$$

$$\gamma = l^l / \eta^{l-1} + o(1/\eta^{l-1}) \text{ as } \eta \rightarrow \infty \quad (41)$$

Figure 2 shows solution to the “steady-state” equation (25) in two cases: $\lambda_1 > \lambda_1^*(\lambda_2)$ and $\lambda_1 < \lambda_1^*(\lambda_2)$, where the “metastability threshold” $\lambda_1^*(\lambda_2)$ for exogenous file transfer load λ_1 monotonically decreases from $\lambda_1^*(0) = 1$ to $\lambda_1^*(\infty) = 0$ with increase in the streaming load λ_2 .




Figure 2. Solution to the equilibrium equation.

If exogenous file transfer load exceeds the metastability threshold: $\lambda_1 > \lambda_1^*(\lambda_2)$, equilibrium equation (25) has no solution $\eta \geq 0$ and, according to the dynamic equation (19)-(20), number of file transfer flows in progress grows infinitely with time for any initial $\eta(0) \geq 0$. If the exogenous file transfer load is below the metastability

threshold: $\lambda_1 < \lambda_1^*(\lambda_2)$, equilibrium equation (25) has two solutions $\eta = \eta^*$ and $\eta = \eta^{**}$ describing respectively stable and unstable equilibriums of the dynamic equations (19)-(20). Number of file transfer flows in progress approaches stable equilibrium if initial number of carried flows is sufficiently small: $\eta(0) < \eta^{**}$. Otherwise, i.e., if $\eta(0) > \eta^{**}$ then the number of file transfer flows in progress grows infinitely with time. Note that thresholds $\lambda_1^*(\lambda_2)$, η^* and η^{**} can easily be evaluated numerically.

More detailed analysis shows that equilibrium η^* represents the desirable metastable network state with finite number of flows in progress. This metastable state can be transformed into the stable state with flow admission control admitting arriving file transfer flows if and only if $\eta < \eta^{**}$. The stabilization is achieved at the cost of small flow rejection probability at the fluid asymptotic regime.

VI. CONCLUSION

This paper discusses a possibility that cascading failures in Internet-type networks with fair bandwidth sharing may be a result of metastability rather than stability of the observed “normal” network state. Sufficient network overload can push the network out of the stability region of the normal state initiating a process of cascading failures. This possibility raises numerous research issues including existence, stability margins, and queuing performance of the metastable state for a general network topology. Most importantly, future research should address the validity of the model and its predictions.

REFERENCES

- [1] J. Mo and J. Walrand, “Fair end-to-end window-based admission control,” *IEEE/ACM Trans. on Networking*, No. 8, pp. 556-567, 2000.
- [2] S. Ben Fredj, T. Bonald, A. Proutiere, G. Regnie, and J. Roberts, “Statistical bandwidth sharing: a study of congestion at flow level,” *SIGCOMM 2001*.
- [3] G. de Veciana, T.J. Lee, and T. Konstantopoulos, “Stability and performance analysis of networks supporting elastic services,” *IEEE/ACM Trans. on Networking*, No. 9, pp. 2-14, 2001.
- [4] P. Key, L. Massoulié, A. Bain, and F.P. Kelly, “Fair Internet traffic integration: network flow models and analysis,” *Annales des Telecommunications* 59, pp. 1338-1352, 2004.
- [5] L. Massoulié and J.W. Roberts, “Arguments in favor of admission control for TCP flows,” *ITC 16*, Edinbourg, June 1999.
- [6] V. Marbukh, “Metastability of fair bandwidth sharing under fluctuating demand and necessity of flow admission control,” *Electronics Letters*, 13th September 2007, Vol 43, No 19.
- [7] R.J. Gibbens, S.K. Sargood, C. Van Eijl, F.P. Kelly, H. Azmoodeh, R.F. Macfadyen, and N.W. Macfadyen, “Fixed-point model for end-to-end performance analysis of IP networks,” *13th ITC Specialist Seminar: IP Traffic Modeling, Measurement and Management*, 2000.
- [8] F. Kelly, “Mathematics unlimited 2001 and beyond,” editors B. Engquist and W. Smit, Springer-Verlag, Berlin, pp. 685-702, 2001.