Metastability of fair bandwidth sharing under fluctuating demand and necessity of flow admission control

V. Marbukh

A flow-level Markov model for fair bandwidth sharing with packet retransmissions and random flow arrivals/departures is proposed and discussed. Despite the model being unstable even under light exogenous load, a desirable metastable state may exist. The network can be stabilised with flow admission control at the cost of small flow blocking probability.

Introduction: A flow level Markov model of fair bandwidth sharing under fluctuating demand has been proposed in [1, 2] for a case of file transfer flows, and in [3] for a case of a mixture of file transfer and streaming flows. These Markov models assume separation of time-scales: given numbers of flows in progress, the fair bandwidth sharing protocol reaches equilibrium much faster than the numbers of flows in progress change due to flow arrivals/departures. Stability under the condition that each link can accommodate its average load was established in [2, 3]. However, these Markov models and stability results do not account for bandwidth wasted on transmissions of 'dead' file transferring packets which will be dropped downstream and then retransmitted.

This Letter proposes to account for wasted bandwidth by assuming that the file transfer rates are determined by the end-to-end goodputs rather than the corresponding throughputs as in [1–3]. The 'goodput-based' Markov model is unstable even under light exogenous load, when the corresponding throughput-based models [1–3] are stable. The instability, observed in [4] by simulations, is a result of demand fluctuations: increase in the number of flows in progress causes increase in packet loss, reducing goodput and further increasing the number of flows in progress. Despite instability, a desirable metastable network state with a finite number of flows in progress may still exist. The network can be stabilised in a close neighbourhood of this metastable state with appropriately designed flow admission control at the price of a small flow rejection probability. Network over provisioning without flow admission control only reduces but does not eliminate the instability region.

Models with fixed set of flows: Consider a network comprised of links $j \in J$ with capacities C_j. Given a vector of numbers of flows $N = (N_r)$ carried on all feasible routes $r \in R$, the fixed point model [5] determines end-to-end flow rates $x = (x_r)$ and packet loss probabilities $p = (p_r)$ for all feasible routes $r \in R$. In a case of small packet losses: $p_r << 1$, (α, w) – fair rate allocation [6] results in the following rate assignments:

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

where α, $w_r > 0$ are some parameters, and end-to-end packet loss on a route $r \in R$ is the sum of packet losses on all the links comprising this route: $p_r = \Sigma p_j$, $j \in r$. Cases $\alpha = 0$, $\alpha = 1$ and $\alpha = \infty$ correspond, respectively, to an allocation which maximises throughput, is proportionally fair or is maxmin fair. Case $\alpha = 2$ describes Transmission Control Protocol (TCP) Reno.

A large number of expressions for source transmission rates in a case of medium and heavy packet losses have been proposed. A straightforward generalisation of (1) is as follows:

$$x_r = \left\{ w_r \Big/ \left[1 - \prod_{j \in r} (1 - p_j) \right] \right\}^{1/\alpha} \left[\prod_{j \in r} (1 - p_j) \right]^{\chi} \tag{2}$$

where $\chi \geq 0$, e.g. [7] considers $\chi = 1/\alpha = 0.5$. When link j carries a large number of flows the following fluid limit approximation can be used:

$$p_j = \max \left(0, 1 - C_j \Big/ \sum_{r: j \in r} N_r x_{rj} \right) \tag{3}$$

where the link j load produced by flows carried on a route r is

$$x_{rj} = x_r \prod_{j \in r_j^-} (1 - p_j) \tag{4}$$

and the part of a route r positioned upstream from link $j \in r$ is r_j^-. Given a vector of numbers of flows in progress $N = (N_r)$, (2)–(4) form a closed system of fixed-point equations for a vector of source rates and link packet losses (x_r, p_j). After solving (2)–(4), the goodput for a flow carried on a route r is determined as follows:

$$g_r = x_r \prod_{j \in r} (1 - p_j) \tag{5}$$

Markov models with arriving/departing flows: Following [3], assume that the network carries file transfer 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 transfer and streaming flows, respectively, carried on all feasible routes $r \in R$. We assume that file transfer and streaming flows arrive on a route $r \in R$ according to a Poisson process of rate Λ_{1r} and Λ_{2r}, respectively. The size of a file arriving on a route $r \in R$ is distributed exponentially with average b_r and the holding time of a streaming flow arriving on a route $r \in R$ is distributed exponentially with average τ_r. All flow arrivals, file sizes and holding times are jointly statistically independent.

We assume separation of time-scales: given vector $N = N_1 + N_2$ of numbers of flows in progress $N = (N_r)$ where $N_r = N_{1r} = N_{2r}$, the flow control protocol reaches equilibrium bandwidth sharing much faster than the number of flows change due to flow arrivals/departures. Under this assumption the evolution of the number of flows in progress can be approximated by a homogeneous in time $t \geq 0$ Markov process $(N_1(t), N_2(t))$. We propose to account for retransmission of file transferring packets by assuming that the file transfer rates are determined by the end-to-end goodputs (5) rather than the corresponding throughputs x_r as in [1–3]. This seemingly minor change drastically alters stability properties of the Markov process $(N_1(t), N_2(t))$ owing to deterioration of the aggregate route goodputs as the numbers of flows in progress increase. Note that separation of time-scales allows for modelling of congestion-dependent admission control and user behaviour by assuming that rates $\Lambda_{ir} = \Lambda_{ir}(N)$ depend on the vector of numbers of flows in progress $N = N_1 + N_2$.

It is known that the ability of links to sustain their average file transfer load:

$$\rho_j \stackrel{def}{=} \frac{1}{C_j} \sum_{r: j \in r} \Lambda_{1r} b_r < 1, \quad \forall j \in J \tag{6}$$

are necessary and sufficient conditions for ergodicity of the 'throughput-based' Markov model when rates Λ_{ir}, $\forall i = 1,2; r \in R$ are fixed [2, 3]. One may expect that the necessary and sufficient stability conditions for the 'goodput-based' Markov model can be obtained from (6) by accounting for the additional load due to retransmissions of file transferring packets:

$$\tilde{\rho}_j \stackrel{def}{=} \frac{1}{C_j} \sum_{r: j \in r} \frac{\Lambda_{1r} b_r}{\prod_{i \in r_j^+} (1 - p_i)} < 1, \quad \forall j \in J \tag{7}$$

where r_j^+ is the part of a route r located downstream from link $j \in r$. Link packet losses p_i increase, and typically approach 1, with increase in the number of flows carried on the link. This may cause instability of the goodputbased model in a case of multihop routes even under light loads, when stability conditions for the throughput-based model (6) are satisfied.

Despite instability, the goodput-based model may have a desirable metastable state with finite number of flows in progress. Under fluid regime

$$\Lambda_{ir} = \varepsilon^{-1} \lambda_{ir},\ C_j = \varepsilon^{-1} c_j,\ \lambda_{ir}, c_j = O(1),\ \varepsilon \to 0 \tag{8}$$

$i = 1, 2$, $\forall r \in R$, $\forall j \in J$, this metastable state corresponds to a locally stable equilibrium (n_{ir}^*), $n_{ir}^* < \infty$ of the following system of ordinary differential equations:

$$\dot{n}_{1r} = \lambda_{1r}(n_1 + n_2) - b_r^{-1} n_{1r} g_r (n_1 + n_2) \tag{9}$$

$$\dot{n}_{2r} = \lambda_{2r}(n_1 + n_2) - n_{2r} \tau_r^{-1} \tag{10}$$

describing evolution of normalised numbers of flows in progress $n_{ir} = \varepsilon N_{ir}$, $i = 1, 2$, $\forall r \in R$ [3]. The metastable state exists under sufficiently light exogenous load and is realised if the initial number of flows in progress is sufficiently small. Otherwise, the number of file transfer flows infinitely grows with time: $\lim_{t \to \infty} n_{1r}(t) = \infty$, $r \in R$, reflecting instability of the corresponding Markov goodput-based model even under light exogenous load λ_{ir}, $i = 1, 2$; $r \in R$, when stability conditions (6) for the corresponding throughput-based model are satisfied. The network can be stabilised in a close neighbourhood of the desirable

Techset Composition Ltd, Salisbury Doc: {EL}Articles/Pagination/EL53520.3d
Telecommunication

metastable state with appropriately designed flow admission strategy at the cost of small flow rejection probability.

Example: Consider a symmetric ring network with K nodes, where each node $k = 1,2,\ldots, K$ is connected to node $(k+1)$ mod (K) by a directed link j_k of capacity C. Let weights $w_r = w$, flow arrival rates $\lambda_{ir} = \lambda_i$, $i = 1, 2$, average file sizes $b_r = b$, average streaming flow durations $\theta_r = \theta$ and numbers of carried flows $n_{ir} = n_i$, $i = 1, 2$ be route $r \in R$ independent, where the set of feasible l-link routes is $R = \{(j_k, j_{(k+1)\mathrm{mod}(K)}, \ldots, j_{(k+l)\ \mathrm{mod}(K)}): k = 1,\ldots,K\}$. We assume that flows arrive only on feasible routes: $\lambda_{ir} = n_{ir} = 0$, $i = 1, 2$, $r \notin R$. Owing to limited space we only consider a case of proportional fair rate assignments (2) with $\alpha = 1$ and $\chi = 0$ under fluid asymptotic regime (8) and assume $w = c$. Introduce utilisation by file transfer flows $\rho = l\lambda_1 b/c$, utilisation by streaming flows $\beta = l\lambda_2\tau$, and normalised numbers of flows carried on a link $\eta_i = n_i l$, $i = 1, 2$. Under all these assumptions the efficiency of the bandwidth sharing, measured by the fraction of link bandwidth occupied by packets to be delivered to their destinations as opposed to packets to be dropped downstream, is

$$\gamma(\eta) = \frac{\eta}{(1+\eta/l)^l - 1} \tag{11}$$

where $\eta = \eta_1 + \eta_2$. Note that (11) follows from (5) and expression for the link packet loss $p = (1 + l/\eta)^{-1}$, which can be derived from (2)–(4).

Consider a case of fixed arrival rates λ_i, $i = 1, 2$ assuming that the numbers of streaming flows in progress already reached equilibrium: $\eta_2 = \beta$ [3]. In this case system (9)–(10) simplifies into the following single differential equation for normalised number of file transfer flows carried on a link:

$$b\dot{\eta}_1 = \rho - \frac{\eta_1}{(1+(\eta_1+\beta)/(l))^l - 1} \tag{12}$$

Equilibriums of (12) are determined by the following fixed-point equation:

$$\eta_1 = f(\eta_1) \overset{def}{=} \rho\left[\left(1+\frac{\eta_1+\beta}{l}\right)^l - 1\right] \tag{13}$$

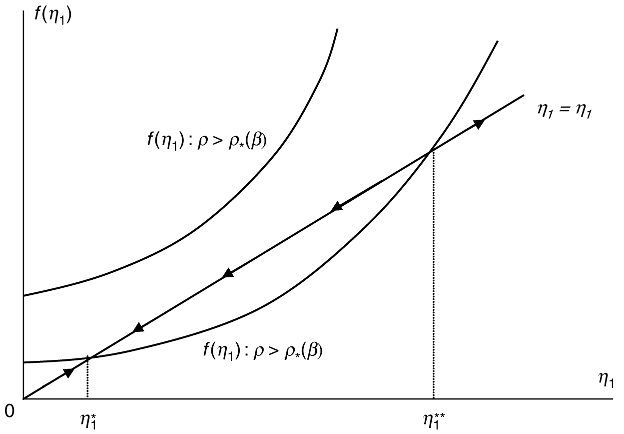

Fig. 1 *Solution to equilibrium fixed point equation* $\eta_1 = f(\eta_1)$

Fig. 1 shows the solution to the 'steady-state' equation (13) in two cases: $\rho > \rho_*(\beta)$ and $\rho < \rho_*(\beta)$, where the metastability threshold for file transfer load $\rho_*(\beta)$ monotonously decreases from $\rho_*(0) = 1$ to $\rho_*(\infty) = 0$ with increase in the exogenous streaming load β. If the exogenous file transfer load ρ exceeds the metastability threshold: $\rho > \rho_*(\beta)$, equilibrium equation (13) has no solution $\eta_1 > 0$, and according to dynamic equation (12) the number of file transfer flows in progress $\eta_1(t)$ infinitely grows with time for any initial $\eta_1(0) \geq 0$. If the exogenous file transfer load ρ is below the metastability threshold: $\rho < \rho_*(\beta)$, equilibrium equation (13) has two solutions $\eta_1 = \eta_1^*(\rho, \beta)$ and $\eta_1 = \eta_1^{**}(\rho, \beta)$ describing, respectively, stable and unstable equilibriums of dynamic equation (12). In this case the number of file transfer flows in progress $\eta_1(t)$ stabilises with time: $\lim_{t\to\infty} \eta_1(t) = \eta_1^*(\rho, \beta)$ if initially the network is lightly loaded $\eta_1(0) < \eta_1^{**}(\rho, \beta)$. Otherwise, i.e. if $\eta_1(0) > \eta_1^{**}(\rho, \beta)$ then the number of file transfer flows in progress $\eta_1(t)$ infinitely grows: $\lim_{t\to\infty} \eta_1(t) = \infty$.

Functions $\rho^*(\beta)$, $\eta_1^*(\rho, \beta)$ and $\eta_1^{**}(\rho, \beta)$ can be easily evaluated numerically. Consider some particular cases. In a case $\beta = 0$, i.e. when the network carries only file transfer flows, the metastability threshold is $\rho^* = 1$. When $\rho < 1$, the metastable state is $\eta_1^* = 0$, and unstable state η_1^{**} is the unique positive solution of equation $\eta/\rho = (1+\eta/l)^l - 1$. This equation yields $\eta_1^{**} = 4(\rho^{-1} - 1)$ if $l = 2$, and takes the form $\eta/\rho = e^\eta - 1$ if $l \to \infty$. When $\beta \to 0$, $\delta \overset{def}{=} 1 - \rho \to 0$, the following asymptotic formulas result from (13):

$$\delta_* \overset{def}{=} 1 - p_* = \sqrt{2\frac{1-1}{l}\beta} \tag{14}$$

$$\eta_1^* = \frac{l}{l-1}\delta - \sqrt{\frac{l}{l-1}\left(\frac{l}{l-1}\delta^2 - 2\beta\right)} \tag{15}$$

$$\eta_1^{**} = \frac{l}{l-1}\delta + \sqrt{\frac{l}{l-1}\left(\frac{l}{l-1}\delta^2 - 2\beta\right)} \tag{16}$$

More detailed analysis shows that equilibrium η_1^* represents the desirable metastable network state with a finite number of flows in progress. This metastable state can be transformed into a stable state with flow admission control, which admits arriving file transfer flows if and only if the normalised number of file transfer flows already in progress η_1 lies within the stability region of this metastable state: $\eta_1 < \eta_1^{**}$. The stabilisation is achieved at the cost of asymptotically small flow rejection probability under the fluid regime.

Conclusion: A flow-level performance model for fair bandwidth sharing with packet retransmissions and arriving/departing flows is proposed. The model accounts for packet retransmissions by assuming that file transfer rates are determined by the end-to-end goodputs rather than the corresponding throughputs as in the conventional model. This seemingly minor modification drastically alters the model stability properties: the goodput-based model is unstable even under light load when the conventional model is stable. Despite instability, a desirable metastable network state with finite number of flows in progress may still exist. The network can be stabilised in a close neighbourhood of the desirable metastable state at the cost of a small flow rejection probability. Future research should address numerous questions raised by the possibility of metastability, such as the effect of the fairness parameters α and w_r in (1)–(2) on the existence, stability margins, and queuing performance of the desirable metastable state as well as a possibility of network stabilisation by a distributed admission strategy for a general topology network.

22 February 2007
Electronics Letters online no: 20070543
doi: 10.1049/el:20070543

V. Marbukh (*National Institute of Standards and Technology, 100 Bureau Drive, Gaithersburg, MD 20899-8920, USA*)

E-mail: marbukh@nist.gov

References

1 Ben Fredj, S., Bonald, T., Proutiere, A., Regnie, G., and Roberts, J.: 'Statistical bandwidth sharing: a study of congestion at flow level'. Proc. of Sigcomm, San Biego, California, USA, 2001

2 de Veciana, G., Lee, T.J., and Konstantopoulos, T.: 'Stability and performance analysis of networks supporting elastic services', *IEEE/ACM Trans. Netw.*, 2001, (9), pp. 2–14

3 Key, P., Massoulie, L., Bain, A., and Kelly, F.P.: 'Fair Internet traffic integration: network flow models and analysis', *Ann. Telecommun.*, 2004, **59**, pp. 1338–1352

4 Massoulie, L., and Roberts, J.W.: 'Arguments in favour of admission control for TCP flows'. ITC 16, Edinburg, UK, June 1999

5 Gibbens, R.J., Sargood, S.K., Van Eijl, C., Kelly, F.P., Azmoodeh, H., Macfadyen, R.F., and Macfadyen, N.W.: 'Fixed-point model for end-to-end performance analysis of IP networks'. 13th ITC Specialist Seminar: IP Traffic Modeling, Measurement and Management, Monterey, USA, 2000

6 Mo, J., and Walrand, J.: 'Fair end-to-end window-based admission control,', *IEEE/ACM Trans. Netw.*, 2000, (8), pp. 556–567

7 Kelly, F.: 'Mathematics unlimited and – 2001 and beyond', *in* Engquist, B., and Schmid, W. (Eds.): (Springer-Verlag, Berlin, 2001), pp. 685–702