



Modified High Frequency Traffic Control Protocol for Congestion Control in TCP Flows

N. NAGENDRA¹, C. GOVARDHAN²

¹PG Scholar, Department of CSE, KMM Institute of Technology & Science,

²Assistant Professor, Department of CSE, KMM Institute of Technology & Science.

Abstract: Congestion control protocols for background data are commonly conceived and designed to emulate low priority traffic, which yields to transmission control protocol (TCP) flows. In the presence of even a few very long TCP flows, this behavior can cause bandwidth starvation, and hence, the accumulation of large numbers of background data flows for prolonged periods of time, which may ultimately have an adverse effect on the download delays of delay-sensitive TCP flows. In this paper, we look at the fundamental problem of designing congestion control protocols for background traffic with the minimum impact on short TCP flows while achieving a certain desired average throughput over time. The corresponding optimal policy under various assumptions on the available information is obtained analytically. We give tight bounds of the distance between TCP-based background transfer protocols and the optimal policy, and identify the range of system parameters for which more sophisticated congestion control makes a noticeable difference. Based on these results, we propose an access control algorithm for systems where control on aggregates of background flows can be exercised, as in file servers. Simulations of simple network topologies suggest that this type of access control performs better than protocols emulating low priority over a wide range of parameters.

Index Terms: Background data, bandwidth sharing, congestion control, delay-sensitive flows.

I. INTRODUCTION

A KEY element of the success of the internet architecture is the ability to accommodate current and future needs of very diverse applications. Connection rates

Differ by many orders of magnitude, while file transfer sizes vary by more than ten orders of magnitude. Nevertheless this is achieved using only a handful of transport protocols, mainly Transmission Control Protocol (TCP) and its variants, which in essence allocate network bandwidth to flows continuously so as to achieve fair sharing at all times. Indeed TCP ‘fairness’ or ‘friendliness’ [1] has become a common prescription for congestion control algorithms which intends to ensure equal sharing between flows.

Network congestion in data networking and queuing theory is the reduced quality of service that occurs when a network node or link is carrying more data than it can handle. Typical effects include queuing delay, packet loss or the blocking of new connections. A consequence of congestion is that an incremental increase in offered load leads either only to a small increase or even a decrease in network throughput. Network protocols that use aggressive retransmissions to compensate for packet loss due to congestion can increase congestion, even after the initial load has been reduced to a level that would not normally have induced network congestion. Such networks exhibit two stable states under the same level of load. The stable state with low throughput is known as congestive collapse. Networks use congestion control and congestion avoidance techniques to try to avoid collapse. These include: exponential backoff in protocols such as CSMA/CA in 802.11 and the

similar CSMA/CD in the original Ethernet, window reduction in TCP, and fair queuing in devices such as routers and network switches. Other techniques that address congestion include priority schemes which transmit some packets with higher priority ahead of others and the explicit allocation of network resources to specific flows through the use of admission control.

A possible solution, violating the end-to-end principle of the internet architecture, is for the internet service providers (ISPs) to intervene and throttle the bandwidth assigned to background data leaving more space for delay-sensitive traffic, or offering some form of prioritization. But this is not in many cases an efficient solution, since the ISPs cannot have the necessary information on how much throttling is necessary, and for which flows [4]. Also, unjustified throttling of traffic can have serious side effects for the ISP business, e.g., legal actions taken by disaffected end users [5]. Recognizing this, internet engineers have developed end-to-end ‘less-than-best-effort’ (LBE) congestion control protocols for background data transfers, such as TCP-LP [6], TCP-nice [7], uTorrent transport protocol [8], LEDBAT [9]. These are typically designed to emulate a low priority transport class which yields to TCP traffic, but this behavior can have a serious drawback under the presence of ‘long’ TCP flows, i.e., persistent or extremely long-lasting and always active flows, as we explain next, motivating our approach. In principle, during the time in which long TCP flows compete with ideal low priority flows, the latter suffer from bandwidth starvation and so their number grows arbitrarily as new low priority flows continue to arrive.

II. RELATED WORK

The subject of fairness between different internet flows is intrinsically linked to congestion control and has been studied extensively over the past decades under different perspectives, e.g., see [13] and references therein. The utility-based approach pioneered in [14] and [15] paved the way for designing new congestion control algorithms for heterogeneous applications [6] and different notions of fairness [17]–[19]. In [10], [18], and [20] the effect of congestion control on the number of ongoing file transfers and download delays is studied. We take a similar viewpoint by considering a model where flow-level dynamics are described by a Monrovia process, and ignore congestion window dynamics and packetlevel effects. Deb et al. [7] consider a flow-level model of a large system with many long and short flows. They consider the optimization of congestion controllers of all flows -background, long and short- by maximizing a social welfare function which includes the average utility obtained by background traffic and the delay caused to short flows. Since we assume that part of the traffic, namely long and short, uses TCP for its transport and cannot be optimized, the optimal policies differ considerably from the ones in [7]. A model where part of the traffic cannot be optimized is considered in [11] where the notion of farsighted congestion controllers for CBF flows is introduced, using a static optimization problem without flow-level dynamics and not involving delays. These controllers implicitly attempt to inflict less delay to short flows but without compromising their average throughput. In this paper we show that these controllers are optimal within the class of policies implemented by state feedback, as they have the same structure as the optimal policy.

III. EXISTING SYSTEM

In Existing System, Congestion control protocols for background data are commonly conceived and designed to emulate low priority traffic, which yields to transmission control protocol (TCP) flows. In the presence of even a few very long TCP flows, this behavior can cause bandwidth starvation, and hence, the accumulation of large numbers of background data flows for prolonged periods of time, which may ultimately have an adverse effect on the download delays of delay-sensitive TCP flows.

IV. PROPOSED SYSTEM

In Proposed System, we look at the fundamental problem of designing congestion control protocols for background traffic with the minimum impact on short TCP flows while achieving a certain desired average throughput over time. The corresponding optimal policy under various assumptions on the available information is obtained analytically. We give tight bounds of the distance between TCP-based background transfer protocols and the optimal policy, and identify the range of system parameters for which more sophisticated congestion control makes a noticeable difference. Based on these results, we propose an access control algorithm for systems where control on aggregates of background flows can be exercised, as in file servers. Simulations of simple network topologies suggest that this type of access control performs better than protocols emulating low priority over a wide range of parameters.

Optimal Sharing Under Dynamic Arrivals of Micro-Flows:

The objective of the model is to minimize the data repositioning costs for data-sharing operators, aiming at a high level users satisfaction, and assuming that it increases with the probability to find an available data at any time. The proposed model considers the dynamic variation of the demand and micro-simulate the DSS in space and time determining the optimal repositioning flows, distribution patterns and time intervals between relocation operations by explicitly considering the route choice for trucks among the stations.

An Access Control Policy for Micro-Flows

we investigate microservice based data processing workflows from a security point of view, i.e., (1) how to constrain data processing workflows with respect to dynamic authorization policies granting or denying access to certain microservice results depending on the flow of the data; (2) how to let multiple microservices contribute to a collective data-driven authorization decision and (3) how to put adequate measures in place such that the data within each individual micro service is protected against illegitimate access from unauthorized users or other micro services. Due to this multifold objective, enforcing access control on the data endpoints to prevent information leakage or preserve one's privacy becomes far more challenging, as authorization policies can have dependencies and decision outcomes cross-cutting data in multiple micro services. To address this challenge, we present and evaluate a workflow-oriented authorization framework that enforces authorization policies in a decentralized manner and where the delegated policy evaluation leverages feature toggles that are managed at runtime by software circuit breakers to secure the distributed data processing workflows. The benefit of our solution is that, on the one hand, authorization policies restrict access to the data endpoints of the micro services, and on the other hand, micro services can safely rely on other data endpoints to collectively evaluate cross-cutting access control decisions without having to rely on a shared storage backend holding all the necessary information for the policy evaluation.

V. CONCLUSION

In this paper a micro-simulation model for optimal relocation of data in data-sharing systems has been presented. In particular, the proposed model jointly determines the optimal carrier data route and the number of data to be repositioned with the relevant stations. The results show that the relocation management increases of users satisfaction in term of probability of finding available data. In particular the proposed DSS is more suitable for non congested bike-sharing systems. The method is modular so that it can be used for wider systems. The proposed LBP reproduce in detailed way the system, thus it can be also used for both strategic design and/or real-time management that is to determine the optimal layout of the LBP. Additional simulations for large-scale LBP with multiple carrier Data are needed and are in progress in order to check the robustness of the method.

REFERENCE

- I. S. Floyd, M. Handley, J. Padhye, and J. Widmer, "Equation-based congestion control for unicast applications," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 30, no. 4, pp. 43–56, Oct. 2000.
- II. F. Dobrian et al., "Understanding the impact of video quality on user engagement," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 41, no. 4, pp. 362–373, Aug. 2011.
- III. L. Schrage, "Letter to the editor—A proof of the optimality of the shortest remaining processing time discipline," *Oper. Res.*, vol. 16, no. 3, pp. 687–690, 1968.
- IV. B. Briscoe, "A fairer, faster Internet," *IEEE Spectr.*, vol. 45, no. 12, pp. 38–43, Dec. 2008.
- V. N. Anderson. (Jul. 2010). Claim Your \$16! Comcast P2P Settlement Now Final. [Online]. Available: <http://arstechnica.com/techpolicy/2010/07/claim-your-16-comcast-p2p-settlement-now-final/>
- VI. A. Kuzmanovic and E. W. Knightly, "TCP-LP: Low-priority service via end-point congestion control," *IEEE/ACM Trans. Netw.*, vol. 14, no. 4, pp. 739–752, Aug. 2006.
- VII. R. Venkataramani, R. Kokku, and M. Dahlin, "TCP Nice: A mechanism for background transfers," in *Proc. 5th Symp. Oper. Syst. Design Implement. (OSDI)*, Boston, MA, USA, Dec. 2002, pp. 329–344.
- VIII. A. Norberg. uTorrent Transport Protocol, accessed on Mar. 15, 2016. [Online]. Available: http://www.bittorrent.org/beps/bep_0029.html
- IX. S. Shalunov, G. Hazel, J. Iyengar, and M. Kuehlewind, Low Extra Delay Background Transport (LEDBAT), document IETF RFC 6817, 2012.
- X. T. Bonald and L. Massoulié, "Impact of fairness on Internet performance," *ACM SIGMETRICS Perform. Eval. Rev.*, vol. 29, no. 1, pp. 82–91, 2001.
- XI. P. Key, L. Massoulié, and M. Vojnovic, "Farsighted users harness network time-diversity," in *Proc. IEEE INFOCOM*, vol. 4. Mar. 2005, pp. 2383–2394.
- XII. C. Courcoubetis and A. Dimakis, "Fair background data transfers of minimal delay impact," in *Proc. IEEE INFOCOM*, Mar. 2012, pp. 1053–1061.
- XIII. J.-Y. Le Boudec, "Rate adaptation, congestion control and fairness: A tutorial," *Ecole Polytechnique Federale de Lausanne, Tech. Rep.*, Dec. 2000.
- XIV. F. Kelly, "Charging and rate control for elastic traffic," *Eur. Trans. Telecommun.*, vol. 8, no. 1, pp. 33–37, Jan./Feb. 1997.
- XV. F. P. Kelly, A. K. Malullo, and D. K. H. Tan, "Rate control in communication networks: