Annual Progress Report:

Self-Similar Network Traffic and Its Control

NSF ANI-9714707

(Year 1: 3/1/98–2/28/99)

Kihong Park (PI)

Department of Computer Sciences
Purdue University
West Lafayette, IN 47907
park@cs.purdue.edu
Tel.: (765) 494–7821
Fax.: (765) 494–0739
www.cs.purdue.edu/people/park

www.cs.purdue.edu/homes/park/nsf-ani-9714707.html
(Project Homepage)

December 16, 1998

1 Overview

The first year (3/1/98–12/14/98) of the project has been very fruitful with significant new results obtained in self-similar traffic control which are described in Section 2. RA funding from the grant has been used to support Tsunyi Tuan—the PI's Ph.D. student—whose thesis topic is on self-similar traffic control, specifically, multiple time scale congestion control for long-range dependent (LRD) traffic. Two M.S. students, Wei Wang and Ambal Balakrishnan, were engaged in implementation aspects of the project involving the transport of real-time MPEG video using packet-level adaptive forward error correction.

Equipment funding has been used to procure four workstations (3 UltraSparc and 1 Intel) and multimedia kits (camera and compression board) for transporting real-time compressed MPEG video. Two of the machines are used as end-to-end sender/receiver stations, one is used as a router, and the fourth is used as a cross traffic source whose traffic is multiplexed at the router. Collectively they realize an experimental platform where the various algorithms could be implemented, tested, and evaluated in working prototype systems.

2 Summary of New Results

The principal results obtained in year 1 revolve around items 1 and 2 in the original proposal—*predictive feedback control* and *adaptive forward error correction*—for self-similar traffic. Predictive feedback control, with the input and feedback from collegues, has been renamed to *multiple time scale congestion control* to more accurately reflect that *explicit* prediction is carried out over large time scales—effective only when long-range dependence is present—and *implicit* prediction is affected by feedback congestion control.

2.1 Multiple Time Scale Congestion Control

2.1.1 Synopsis

The goal of multiple time scale congestion control (MTSC) is to exploit correlation structure present in long-range dependent traffic at larger time scales (e.g., 2–5s) using *explicit* prediction to affect the behavior of feedback congestion controls acting at smaller time scales (e.g., 20-200ms) which are limited to *implicit* prediction afforded by feedback, to ultimately affect improvement in performance. This relationship is shown in Figure 2.1. Two tasks need to be addressed.

- *On-Line Prediction* A basic requirement of multiple time scale congestion control is on-line prediction of correlation structure in LRD traffic for subsequent exploitation in traffic control. We have shown that LRD traffic with Hurst parameter in the range 0.75–1 possesses strong correlation structure in the 1–10s range that can be effectively detected on-line.

  We employ a maximum likelihood estimator that is amenable to efficient on-line implementation to compute conditional probabilities of the form $\Pr\{ L_2 \mid L_1 = \ell \}$ where L_1 is the traffic level observed over the recent past $[t - T_1, t)$ and L_2 is the predicted traffic level over the near future $[t, t + T_2)$. Optimal prediction has a long history and is not the subject matter of our investigation. Our prediction scheme is *sufficient* for traffic control purposes but any other prediction scheme can be substituted without affecting the conclusions.

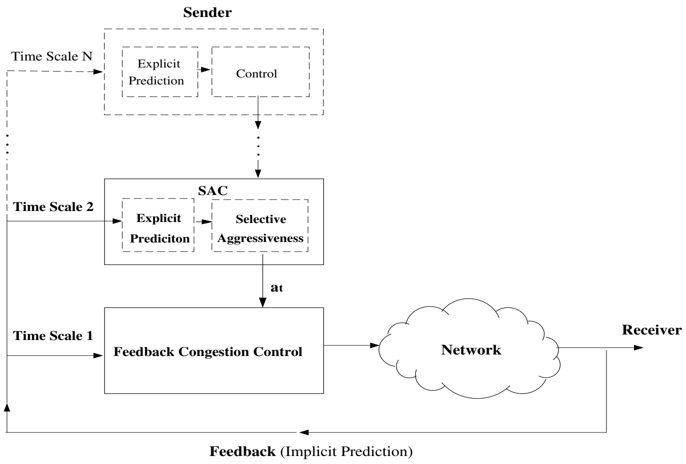

Figure 2.1: Multiple-time scale congestion control framework. Information at Time Scale 2 is used to influence the behavior of congestion control at Time Scale 1. Dashed lines show the potential extensibility of the framework to three or more time scales.

- *Coupling* Coupling refers to the effective exploitation of predicted traffic levels at larger time scales for use in the control actions at smaller time scales. There are a multitude of approaches to coupling congestion control modules acting at different (i.e., in our case two) time scales such that in cooperation and unison they affect cohesive changes to the data rate such that performance is optimized.

  We have developed an approach called *Selective Aggressiveness Control* (SAC) which modulates the aggressiveness exhibited by the smaller time scale feedback congestion control module such that overall aggressiveness of bandwidth consumption is amplified when the larger time scale network contention level is deemed low—i.e., available bandwidth is high—and bandwidth consumption behavior is rendered conservative if the contention level is deemed high. SAC's *raison d'etre*, in turn, is intimately tied to the intrinsic difficulty of achieving optimum throughput in congestion control while maintaining stability.

Finding effective couplings and architecting their efficient realization in new and existing protocols is the central problem that needs to be solved in the context of multiple time scale congestion control. SAC achieves coupling through an aggressiveness parameter a_t which is exported by the large time scale module to the small time scale module which then influences the latter's data rate adjustment behavior. For example, in the case of a generic linear increase/exponential decrease rate based feedback congestion control

$$\frac{d\lambda}{dt} = \begin{cases} \delta, & \text{if } d\gamma/d\lambda > 0, \\ -\varepsilon\lambda, & \text{if } d\gamma/d\lambda < 0, \end{cases} \tag{2.1}$$

where γ denotes throughput, the scaling factor ε of the linear increase phase is equated to a_t, i.e., $\varepsilon \equiv a_t$. Most feedback congestion control algorithms have such parameters that can be modulated by the large time scale congestion control module which would be otherwise held fixed. The control problem is nontrivial since increasing a_t, when the contention level is high, can result in a decrease in performance, and vice versa.

2.1.2 Main Accomplishments

Our main accomplishments for MTSC have been in the identification, formulation, and demonstration of the feasibility of the multiple time congestion control framework for affecting improved traffic control under self-similar traffic conditions [4, 8]. More recently, we have shown the applicability and extensibility of the MTSC approach to self-similar traffic control by implementing SAC on top of TCP with significant performance gains.

Rate-Based Feedback Congestion Control The proof-of-concept of the MTSC approach was demonstrated in [11, 12] where we showed that performance of a generic rate-based feedback congestion control such as (2.1) can be improved significantly (40% and higher) under self-similar traffic conditions by exercising multiple time scale congestion control via SAC. We have shown that throughput gains are highest when long-range dependence is strongest (i.e., $H \nearrow 1$) but SAC does not negatively impact performance when traffic correlation is only short-range.

We also showed that *perfect* prediction—i.e., a priori knowing the future network state with certainty—only yields incremental performance improvement vis-à-vis performance based on *on-line* prediction, and bandwidth is equitably shared as the number of SAC-based MTSC connections sharing a common bottleneck link is increased.

TCP-Based Implementation More recently [13, 10], we have shown the applicability and extensibility of MTSC by implementing SAC on top of TCP (i.e., TCP Reno and several performance-enhanced variants) and showing performance gains in the 40–50% range even when TCP Reno's performance was further enhanced (up to 30%) through separate means not involving MTSC. In the case of TCP Reno—without rate-based enhancements such as those of TCP Vegas—the aggressiveness parameter a_t of SAC was imported by replacing the congestion window update procedure of the linear increase phase `CongestionWindow` $\leftarrow$ `MSS` $\cdot$ `MSS/CongestionWindow` by `CongestionWindow` $\leftarrow a_t \cdot$ `MSS` $\cdot$ `MSS/CongestionWindow`. A similar coupling was affected in the performance-enhanced variants.

In essence, SAC, by controlling the aggressiveness of TCP during its linear increase phase as a function of predicted network state, is able to facilitate aggressive consumption of available bandwidth when the contention level is low, and affect conservative bandwidth consumption when the contention level is high.

2.1.3 Broader Impact

MTSC has the potential of impacting the design of next generation congestion control protocols due to its promising prowess, generality, and flexibility at adapting to existing feedback congestion control protocols. Multiple time scale congestion control can be viewed as imparting *proactivity* to feedback congestion controls above and beyond what can be affected through techniques based on *implicit prediction* alone as afforded by feedback. Indeed, we show that after improving the performance of TCP Reno by up to 30% via implicit prediction techniques we are still able to further improve performance up to 50% using SAC-based MTSC.

One of the principal applications of MTSC is to congestion control in high-bandwidth (B-ISDN) wide area networks where the limitations of feedback congestion control in high *delay-bandwidth product* networks loom as a potentially debilitating problem. Even when deviating from the end-to-end paradigm and architecting control mechanisms that actively engage the internals of networks—i.e., routers—to affect a cascade of feedback loops (e.g., an extreme example being H.T. Kung's credit-based flow control), the problem of high delay-bandwidth product gets reduced but it does not go away. MTSC provides a robust mechanism for exercising proactivity under self-similar traffic conditions[1] when bandwidth and RTT are high. We expect to contribute to IETF's standards bodies in year 2 by proposing MTSC in the form of SAC and TCP WeLa as future RFCs.

2.2 Adaptive Forward Error Correction

2.2.1 Synopsis

A second, important thrust of our work in self-similar traffic control is packet-level adaptive forward error correction for managing the effect of burstiness when satisfying real-time constraints and facilitating QoS. In wide area networks, it is, in general, infeasible to use ARQ-based retransmission schemes to handle dropped packets for real-time traffic streams since the presence of large RTTs can render retransmitted information outdated and thus useless.

Adaptive forward error correction (AFEC) injects redundancy into a real-time transmission through packet-level FEC as a function of network state such that redundancy is increased when a network becomes lossy to compensate for dropped packets. On the other hand, redundancy is decreased when the network is well-behaved to minimize unnecessary bandwidth consumption. The control problem is nontrivial due to the fact that increasing redundancy beyond a certain point can cause self-induced congestion and thus decrease the probability of timely recovery of information at the receiver. From a queueing perspective, packet-level FEC induces correlation structure in the drop or erasure process that stems from queueing effects—especially queueing under self-similar bursty input—leading to dependency structure that is different from those covered previously for various channels with dependencies in bit-level FEC.

Ultimately the goal of AFEC is to facilitate an invariant target end-to-end QoS that is immune to network effects, modulo unpredictable nonstationary changes in network state which affect a transient degradation in QoS due to adaptability.

2.2.2 Main Accomplishments

The main accomplishments of the adaptive FEC project are providing a comprehensive analysis of the adaptive redundancy control problem for QoS-sensitive transport of real-time traffic and implementation of AFEC for transport of real-time MPEG video. Self-similar burstiness is introduced by the nontrivial correlation structure exhibited by compressed VBR video and explicit LRD cross traffic that is multiplexed with AFEC traffic at a shared router.

Stability Analysis In [5], we analyze the optimal redundancy control problem associated with adaptive FEC and show conditions for stability. We devise a specific protocol called AFEC for QoS-sensitive transport of real-time traffic and show its effectiveness by analysis and simulation. AFEC has the property of being a *symmetric* control—the upward and downward rate adjustments are reflections of each other—in the stable regime, whereas in the unstable regime, it reverts to *asymmetric* control where downward adjustment is exponential. The *target* end-to-end QoS is specified by the end user (i.e., receiver), and thus, if the desired operating point falls into the stable

[1] We remark that our results are impervious to the debate going on in some quarters about whether network traffic is technically long-range dependent or not since our performance results are based on *measured* traffic which do not depend on exact mathematical explication of correlation structure except that it exists over a *finite* range of time scales relevant to congestion control.

regime, adaptive control is affected in a smooth manner by symmetric control facilitated by the asymptotic stability of the target operating point. Only when the target QoS falls outside of the stable regime—i.e., network resources are insufficient to facilitate the desired QoS—is asymmetric control invoked.

More recently [9], we have extended the stability analysis of adaptive FEC to a 2-dimensional system which allows us to draw more refined conclusions regarding its behavior, especially, in the context of receiver-oriented QoS control where symmetric and asymmetric control coexist.

Real-Time MPEG Video Transport We have implemented a version of AFEC—called MAFEC (AFEC for MPEG video)—for the transport of real-time MPEG video under Solaris UNIX running on UltraSparc workstations. The receiver specifies a parameter $0 \leq \gamma \leq 1$—in addition to several other parameters including frame rate f—which represents the user's desired QoS where γ represents the fraction of video frames that are received and rendered in a timely manner. Thus $\gamma = 1$ means that all frames are to be delivered and processed in a timely manner commensurate with the given frame rate. It is MAFEC's job to satisfy the user-specified target QoS assuming network resources are adequate for the task at hand. The overall structure of the system is depicted in Figure 2.2 (left).

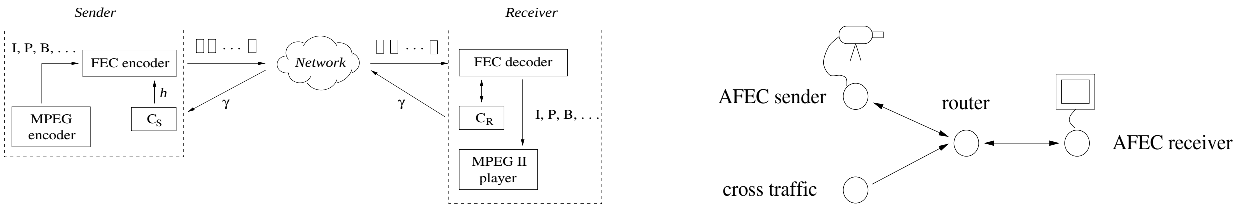

Figure 2.2: Left: Overall structure of MAFEC. At the sender, I, P, B frames forwarded from an MPEG encoder are fed into the FEC encoder. The latter is controlled by a sender-side controller C_S. At the receiver, the coded packets are fed into the FEC decoder which is controlled by a receiver-side controller C_R. The decoded I, P, B frames are forwarded to an MPEG-II player. Right: Experimental set-up for MAFEC performance measurements.

QoS control is achieved completely end-to-end and the only system dependence is due to real-time video capture/compression and rendering. In most recent work, we are in the process of migrating to a capture/compression platform based on Optibase's Moviemaker real-time compression board for MPEG I video/audio which runs on top of Windows NT (or 98/95). We have used the experimental set-up shown in Figure 2.2 (right) to test and benchmark the performance of MAFEC with promising results [7]. Some initial results are reported in [6].

Real-Time MPEG Audio Transport Most recently, we have begun work on extending our work with MAFEC to MPEG audio. An important subtask of this work is the characterization of compressed VBR audio as a time series. Preliminary results indicate that there is significant burstiness and correlation structure associated with MPEG I audio, however, a more definitive conclusion needs to await the completion of on-going analysis. Interestingly, consultation with several collegues has revealed that comprehensive traffic analysis of compressed VBR audio may not have been previously carried out—as has been done extensively with compressed VBR video—perhaps, due to the comparatively low bit rate associated with audio even in the case of CD quality. We have implemented a prototype system of AFEC customized to MPEG audio with enouraging but yet preliminary results [1]. We interface with the `mpg123` player to render audio on the receiver station.

2.2.3 Broader Impact

We expect the adaptive FEC work to influence the design of protocols for *end-to-end* QoS control in other contexts due to its generality with respect to actively controlling redundancy—a form of proactivity—to facilitate invariant end-to-end QoS over a black box network. For supporting real-time traffic over wide area networks, the only two solutions are resource reservation coupled with admission control and adaptive FEC. The former is able to provide *guaranteed* services—i.e., protection via GPS packet scheduling and desired QoS via resource reservation and admission control—however, its cost in terms of resource allocation is commensurately high.

With the exploration of differentiated services (e.g., IETF's diff-serv Working Group) including the PI's work in Stratified Best-Effort Service (SBS) architecture [2, 3] which provide graded services, albeit, with weaker protection than guaranteed services, *end-to-end* QoS control in the form of AFEC and its spin-offs may aid in imparting further robustness to the rendered services and their QoS. Viewed from another perspective, although SBS, Clark's Assured service, and Jacobson's Premium service all require some router support, this support has been kept to a minimum with end-to-end control at the edges being integral components. AFEC-based end-to-end QoS control adds a further layer of control to the end system consistent with the spirit of a weaker form of end-to-end paradigm.

3 Planned Work for Years 2 & 3

The planned work for years 2 & 3 is consistent with the tasks described in the original proposal and is not repeated here for brevity. A number of new and synergistics tasks—in addition to those described in the original proposal—have sprung up and they have been described in Section 2 in conjunction with already accomplished tasks. The thrust of the work remains unchanged and is consistent with the original proposal.

One synergistic task that we have not mentioned is the use of multiple time scale congestion control in conjunction with AFEC in the transport of real-time traffic. AFEC is proactive in the sense that it actively exploits redundancy to hide various network effects such as packet drops and delays, however, it still relies on feedback control to adjust the level of redundancy as a function of implicitly predicted network state. MTSC has the potential of further improving the performance of AFEC under long-range dependent traffic conditions by reducing the occurrence of premature decrease and increase in redundancy based on nonpersistent short-term effects. Responsiveness to short-range correlations is conducive to minimizing unnecessary bandwidth consumption, however, this comes at the price of incurring adaptation costs—which can translate to reduced QoS—to short-term fluctuations that, by definition, are difficult to predict and transient in nature. MTSC, with AFEC acting as the low time scale module, may enhance QoS by affecting changes in redundancy when the underlying network state changes are persistent.

4 Related Publications

[1] A. Balakrishnan and K. Park. End-to-end transport of real-time MPEG audio using adaptive FEC. Manuscript in preparation, 1998.

[2] S. Chen and K. Park. A distributed protocol for multi-class QoS provision in noncooperative many-switch systems. In *Proc. IEEE International Conference on Network Protocols*, pages 98–107, 1998.

[3] S. Chen and K. Park. An architecture for noncooperative QoS provision in many-switch systems. To appear in *Proc. IEEE INFOCOM '99*, 1999.

[4] K. Park, G. Kim, and M. Crovella. The protocol stack and its modulation effect on self-similar traffic. In K. Park and W. Willinger, editors, *Self-Similar Network Traffic and Performance Evaluation*. Wiley Interscience, 1999.

[5] K. Park and W. Wang. AFEC: an adaptive forward error correction protocol for end-to-end transport of real-time traffic. In *Proc. IEEE International Conference on Computer Communications and Networks*, pages 196–205, 1998.

[6] K. Park and W. Wang. AFEC: an adaptive forward error correction protocol for end-to-end transport of real-time traffic. Technical Report CSD-TR-98-025, Department of Computer Sciences, Purdue University, 1998.

[7] K. Park and W. Wang. QoS-sensitive transport of real-time MPEG video using adaptive forward error correction. To appear in *Proc. IEEE Multimedia Systems '99*, 1999.

[8] K. Park and W. Willinger, editors. *Self-Similar Network Traffic and Performance Evaluation*. To be published by Wiley Interscience, 1999.

[9] Kihong Park. Receiver-oriented QoS control: Symmetry and stability. Manuscript in preparation, 1998.

[10] Kihong Park. Congestion control for self-similar network traffic. To be presented at *European Institute for Probability, Statistics and Stochastic Operations Research,* Workshop on Heavy Tails and Queues, 1999.

[11] T. Tuan and K. Park. Multiple time scale congestion control for self-similar network traffic. Submitted for publication, 1998.

[12] T. Tuan and K. Park. Congestion control for self-similar network traffic. In K. Park and W. Willinger, editors, *Self-Similar Network Traffic and Performance Evaluation*. Wiley Interscience, 1999.

[13] T. Tuan and K. Park. TCP WeLa: Achieving proactivity through multiple time scale congestion control. Submitted for publication, 1999.