Politecnico di Bari
Dipartimento di
Ingegneria Elettrica
e dell'Informazione

C3LAB
Control of Computing
and Communication
Systems Lab

A Hybrid Model of the Akamai Adaptive Streaming Control System

Cape Town, South Africa

26 August 2014

L. De Cicco, **G. Cofano** and S. Mascolo

Politecnico di Bari, Dipartimento di Ingegneria Elettrica e dell'Informazione

1. Introduction
2. The Akamai Control System
3. Model and Properties
4. Experimental Validation
5. Conclusion

INTRODUCTION

- In video streaming a client streams the video over an HTTP connection from the server that stores it;
- Video streaming is becoming the largest fraction of the Internet traffic;
- With the stream-switching approach the video bitrate can be throttled on-the-fly to match the time-varying available bandwidth;
- At the client a buffer is employed to absorb instantaneous bandwidth variations;

THE STREAM-SWITCHING APPROACH

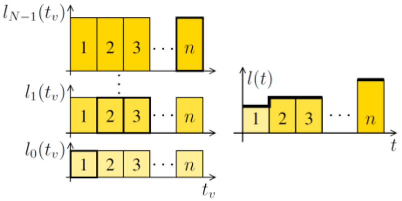

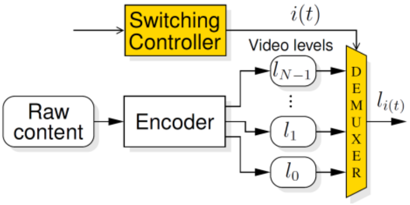

- The raw video is encoded in N different qualities (video levels or bitrates) which are stored at the server;
- Each level is divided into segments of fixed duration;
- At every segment download the Switching Controller selects the level of the next segment;

THE PLAYOUT BUFFER MODEL

In a generic video streaming control system the playout buffer length, i.e. the total duration of video stored in the playout buffer, can be modelled as:

$$\dot{q}(t) = f(t) - d(t)$$

filling rate:

$$f(t) = \frac{r(t)}{l(t)}$$

r(t) is the received rate, l(t) is the selected video bitrate

draining rate:

$$d(t) = \begin{cases} 1, & playing \\ 0, & paused \end{cases}$$

THE AKAMAI STREAM-SWITCHING CONTROL SYSTEM

Motivation: why are we considering the Akamai stream-switching control system?

- it is a leading CDN operator whose platform is employed by several streaming platforms, including Livestream;
- it presents some interesting and perhaps unique control features;

… however, the code is not open source, thus this work is based on a previous work of system identification!

(see De Cicco, L. and Mascolo, S.,“*An adaptive video streaming control system: Modeling, validation, and performance evaluation”*. IEEE/ACM Transactions on Networking)…

THE AKAMAI STREAM-SWITCHING CONTROL SYSTEM

The Akamai control system consists of *two controllers*:

1. the *stream-switching controller*, which selects the video level;
2. the *playout buffer length controller*, which aims at avoiding that the buffer gets empty (buffering events) by throttling the sending rate;

Their behavior changes according to the logical phase in which the control system is. There are *three phases*:

1. the *Buffering* phase: entered to quickly fill the queue at the start or after a buffering event;
2. the *Normal* phase (periodically triggered): the sending rate is throttled to steer the buffer length to a target length;
3. the *Greedy* phase (periodically triggered): the sending rate is set to a much higher value than in the *Normal* phase to probe and estimate the available bandwidth;

THE AKAMAI STREAM-SWITCHING CONTROLLER

The *stream-switching controller*:

- given the available bandwidth, selects the optimal video level $l_{opt}(t)$ such that:

$$l_{opt}(t) = \underset{l \in L}{argmax}\ l$$
$$s.t.\ \hat{b}(t) > (1 + S_f)l$$

estimated bandwidth ($\hat{b}(t)$)

Safety margin $S_f > 0$

- is event-based:
    1. level switch-up is triggered when $\hat{b}(t) > (1 + S_f)l_{i+1}$
    2. level switch-down is triggered when $q(t) < q_d$

lower threshold (q_d)

THE AKAMAI PLAYOUT BUFFER LENGTH CONTROLLER

The *playout buffer length controller* throttles the sending rate to steer the buffer length to the target q_T:

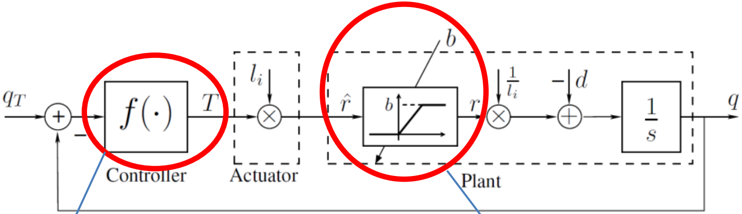

T(t) is equal to:

1. $T_B > 1$ (*Buffering* phase);
2. $\max(1 + \frac{q_T - q(t)}{q_T}, T_m)$ (*Normal* phase);
3. $T_M > 2$ (*Greedy* phase);

the saturation block models the bottleneck link of bandwidth b(t)

MODELLING ASSUMPTIONS

- A piecewise constant bandwidth input function has been employed. It allows us to analyze any practical traffic scenario with a bottleneck link;
- The communication forward and backward delays from the client to the server and the actuation time-delay have been neglected;
- The safety factor S_f, the queue threshold q_d and the queue target q_T have been assumed to be constant (in the reality they are time-varying);

THE HYBRID AUTOMATON

Due to the state-dependent and event-triggered dynamics and the discontinuous elements the model has the form of the following hybrid automaton:

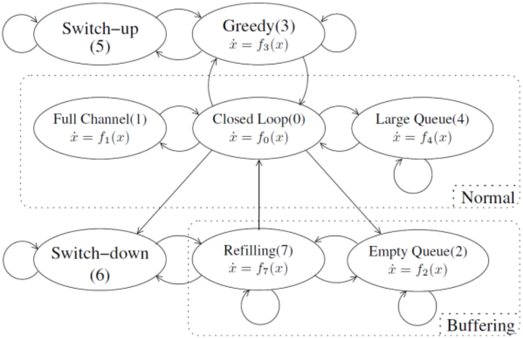

The state $x = [i, r, q, \tau, \tau_1, \tau_2]^T$:

- i (the video level index);
- r (the sending rate);
- q (the queue length);
- τ (the phases timer),
- τ_1 (time-varying duration of the Normal phase timer);
- τ_2 (time−varying duration of the Greedy phase timer);

PROPERTIES

Proposition 1: a necessary condition to permit a switch-up between two adjacent levels l_i and l_{i+1} is that $\frac{l_{i+1}-l_i}{l_i} \leq \frac{T_M}{S_f}$.

Remark: this condition ensures the reachability of all the video levels.

Proposition 2: a sufficient condition for the boundedness of the playout buffer length $q(t)$ is that $\frac{\tau_2}{\tau_1} \leq \frac{1-T_m}{T_M-1}$.

Remark: large buffering is a waste of resources from the network point of view. It is prevented when this condition is satisfied.

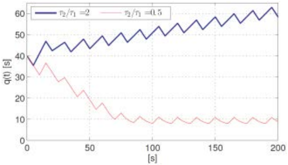

Proposition 3: let us assume that an actuation time-delay τ_a occurs when a switch-down event is triggered and that $B \geq l_0$. A sufficient condition to avoid buffering is that $q_d > (1 - \frac{l_0}{l_i})\tau_a$.

Remark: thanks to this property the threshold q_d can be tuned to ensure robustness against actuation time-delays.

THE MODEL VALIDATION PROCEDURE

- the dynamics of the variables of the simulated model have been compared with the ones obtained through Internet experiments;

- the model has been simulated with the Matlab *Hybrid Equations (HyEq) Toolbox;*

- the experimental results have been obtained by playing the video “Elephant’s Dream”, served by the Akamai server, on a Linux PC equipped with the traffic shaper **tc** to change the link capacity in real time;

- abrupt step-like bandwidth increases and decreases have been considered to validate, respectively, the switch-up and the switch-down case;

THE SWITCH-UP CASE

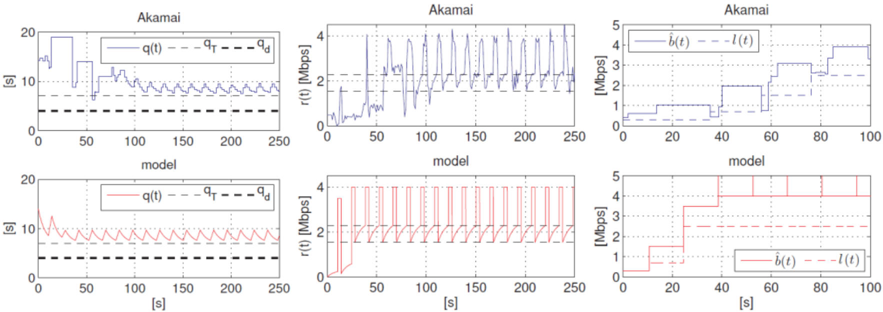

- The model is quite accurate;
- The succession of the two phases is confirmed;
- The difference in the transient is due to the neglected actuation delay, which is present in the real system;

THE SWITCH-DOWN CASE

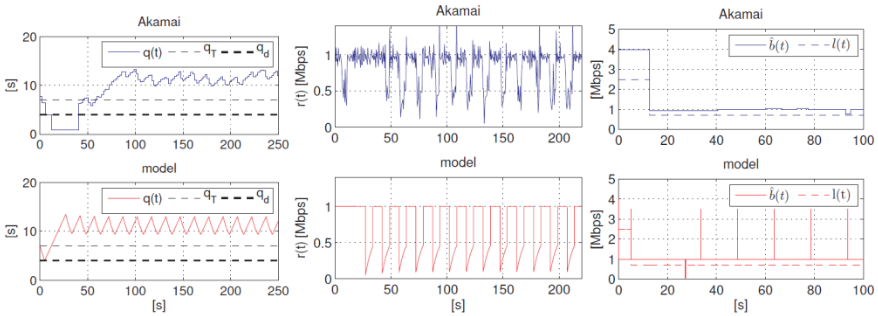

- The model is quite accurate;
- The alternating dynamics due to phases are shown;
- Here, too, the difference in the transient is due to the neglected actuation delay;

CONCLUSIONS

- A model of the Akamai stream-switching control system has been proposed;
- The model is in the form of a hybrid automaton;
- The model has been validated by comparing simulations and experimental results;
- We have provided insights on the parameters tuning by means of some key properties;

Questions?