AVAILABLE BIT RATE TELETRAFFIC CONTROL USING MODEL PREDICTIVE CONTROL

Robert R. Bitmead

Department of Mechanical & Aerospace Engineering, University of California San Diego
9500 Gilman Drive, La Jolla CA 92093-0411, USA
Email:`rbitmead@ucsd.edu` Fax: +1 858 822 3107

ABSTRACT

Model Predictive Control is analyzed as a strategy for coping with the unknown, variable and non-deterministic action and measurement delays which can occur in congestion control of available-bit-rate teletraffic control for high-speed data networks. It is shown that the state-feedback-plus-estimator certainty equivalence nature of this control law is particularly amenable to coping with these delays. Both the estimator update rule and the receding-horizon controller can be modified to accommodate absent network measurements. This is standard for the estimator. But involves a novel approach to the receding-horizon control law, which for the congested network results in diminishing the network utilization objective in favor of queue stability.

1. INTRODUCTION

Teletraffic in high-speed ATM networks consists of a number of bit-rate regimes depending on the quality-of-service (QoS) guarantee agreed between the sender and the network service provider. *Constant bit-rate* (CBR) traffic and *variable bit-rate* (VBR) are guaranteed QoS options which have high-priority bit-rate assignment in response to demand. *Available bit-rate* (ABR) traffic, by contrast, is a best efforts service in which traffic is assigned bandwidth in response to its availability, given the demands of the CBR and VBR sources and the appearance and disappearance of other ABR traffic. Since the bit-rates assigned to ABR traffic may vary with availability, ABR represents the only part of the traffic which is amenable to feedback control in response to network congestion.

Research supported by USA National Science Foundation Grant ECS-0070146 and DARPA Grant N00014-00-1-0799

The control objective here is twofold; to maximize network utilization or traffic throughput, and to minimize the rate of lost packets and consequent retransmission inefficiencies. The feedback control is applied by communicating the maximum allowable transmission bit-rate to each ABR source in response to *resource management* (RM) packets, which are interleaved with the data traffic in ATM networks and convey congestion information. See [1, 2, 3] for recent surveys and analyses of the formulation of this stochastic control problem. The stochastic nature flows from the need to include a model for the variation of ABR bandwidth — an autoregressive model is used in [1] for example. Apart from this stochastic variability, the dominant feature of the ABR congestion control problem is the presence of a non-negligible delay in the feedback path linking the bottleneck constrained node's notification of capacity to the ultimate control of the source ABR data rate. The *action delay* or *round trip time* is assumed both fixed and known in the above references. Mascolo [2] uses the Smith Predictor to develop feedback controllers capable of yielding stability with this delay. Altman *et al* [1] generate *certainty equivalence* controllers based on optimal state-variable feedback combined with state estimation with delayed measurements and again establish stability for these schemes.

Our approach here is to study the suitability of Model Predictive Control (MPC) approaches for ABR stability and performance control. In particular, we view MPC as a variant of certainty equivalence control and show how it is suited to coping with variable and unknown action delays in this circumstance. We use the fact that MPC is comprised of a state estimator and a receding-horizon optimal state-feedback control law to develop separate delay-management modifications to each part. The motivation for considering this formulation of variable and *a priori* unknown delay is that congestion in ATM networks leads to reduced ABR bandwidth, which in turn leads to dynamically decreased bit-rate assignment and increased communication delay. The asymmetry of traffic flows in each direction also can cause variations in round-trip delay between source and destination or intermediate nodes.

2. NETWORK CONGESTION CONTROL MODEL

We consider a store-and-forward packet switched network, one node of which is depicted in Figure 1. We follow the

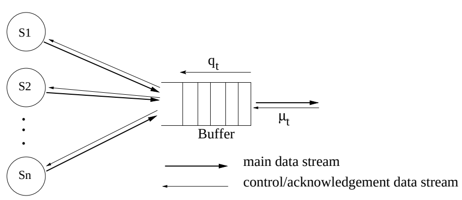

Fig. 1. Network node with ABR sources.

formulation and notation of [1]. The queue length at the node is denoted by q_t and has a maximal value of Q. Packets arriving after the queue is full are lost and must be retransmitted. The available downstream or service bit-rate is denoted by μ_t, which is a stochastic process reflecting available network resources. The source data rate at the node due to source S_i is $r_{i,t}$. Note that, because of transmission delays, this source rate is due to data injected a number of time steps earlier. The total time for a source rate decision from the node to reach the source and for that data rate of arrivals to reach the node is d_i. The commanded rate from the source, denoted $v_{i,t}$, is thus given by

$$v_{i,t-d_i} = r_{i,t}.$$

The queue length at the node evolves according to

$$q_{t+1} = q_t + \sum_{i=1}^{n} r_{i,t} - \mu_t = q_t - \sum_{i=1}^{n} v_{i,t-d_i} - \mu_t. \quad (1)$$

In this framework, the specification of commanded rate, $v_{i,t}$ is the control signal to the sources and variation in μ_t is the stochastic disturbance to the measured process output, q_t. The action delays, d_i, plus the fact that downstream bottlenecked nodes need to communicate their constraints make this a difficult control problem,.

Stability of the system corresponds to the property that

$$q_t \leq Q, \quad \text{for all} \;\; t > 0. \quad (2)$$

That is, stability corresponds to the absence of queue overflows and consequent lost packets and retransmissions. (In practice, some small component of retransmission might be acceptable.) Clearly, this might be achieved by setting $v_{i,t} = 0$. Therefore the competing network utilization or *performance* objective is stated as

$$q_t > 0, \quad \text{for all} \;\; t > 0. \quad (3)$$

Such an objective countermands the stabilization solution of setting all ABR source rates to zero.

The state description (1) is a linear system, which, when combined with a similar model for the stochastic disturbance μ_t as in [1], yields an expression of the control problem as linear feedback control of a system with delay. This is posed as a Smith Predictor in [2] and as LQG control in [1]. In both cases the delays are assumed fixed and known *a priori*. In practice, the delays are unknown, time-varying and non-deterministic, since they themselves depend on prevailing network traffic.

We study in this paper the applicability of Model Predictive Control (MPC) for the stabilization and control of network traffic in which the variable delay in arrival of measurements from distant parts of the network needs to be accommodated. An approach to dealing with unavailable measurements is simply stated.

2.1. MODEL PREDICTIVE CONTROL

For a general system with inputs u_t, outputs y_t and state x_t, MPC is composed as follows.

Step 1: At current time t, a finite-horizon (N), open-loop, constrained optimal control problem is posed from the current state estimate, $\hat{x}_{t|t-1}$.

$$\{u_t^t, u_{t+1}^t, \ldots, u_{t+N-1}^t\} = \arg \min_{\{u_t^t, u_{t+1}^t, \ldots, u_{t+N-1}^t\}} \left[J(\hat{x}_{t|t}, N) \right], \quad (4)$$

where the optimization is given by

$$\begin{aligned} J(x_t, N) \quad &= x_{t+N}^T S_N x_{t+N} \\ &+ \sum_{i=0}^{N-1} (L(x_{t+i}, u_{t+i})), \quad (5) \\ \text{subject to} \quad & x(t+i+1) = f(x_{t+i}, u_{t+i}) \\ u_{t+i} \in \mathcal{U}_i, \quad & x_{t+i+1} \in \mathcal{X}_{i+1}, \; i = 0, N-1, \end{aligned}$$

where $\mathcal{U}_i$ are the sets of admissible controls and $\mathcal{X}_i$ are the sets of allowable state values.

Step 2: Of $\{u_t^t, u_{t+1}^t, \ldots, u_{t+N-1}^t\}$, the time-$t$ finite-horizon solution sequence, only the first value u_t^t is applied as a control input.

Step 3: A system measurement $y_t = h(x_t, u_t) + v_t$ is taken (perhaps corrupted by noise v_t) and the next state estimate $\hat{x}_{t+1|t}$ computed.

This loop is completed at each time t. This defines a *Receding Horizon* control strategy because, although an N-step-ahead finite horizon control is computed, only the first element in the solution sequence is applied as a control before a new N-step sequence is computed with a sliding horizon.

Much has been written about MPC recently, [4, 5, 6]. A very useful survey is given in [7]. Traditionally, the focus in MPC has been on the properties of the full-state, receding horizon control law. The connection to a state estimator has been somewhat overlooked. There are many appealing features of MPC, which has its historical genesis in Chemical Process Control.

- The finite-horizon, constrained, optimal control problem has an open-loop solution. This therefore admits the solution via explicit criterion minimization, such as might be achieved via direct gradient methods [8].
- Feedback is effected through the state estimate — a factor made explicit in our notation. The receding horizon control's dependence just on the most recent state estimate means that feedback control is implemented.
- The criterion minimization is based on the outputs of a simulation model. This tight coupling between a model, which is frequently available but is not usually amenable to direct controller design, and an ensuing controller achieved implicitly through the minimization process is a major drawcard of the method.
- While MPC proceeds via successive finite-horizon solutions, it is an infinite-horizon control strategy, since the finite horizon problem is re-posed and re-solved at each instant of time. Accordingly, one may pose asymptotic stability questions for MPC. The most powerfully general result stems from Keerthi and Gilbert [9]. Subject to other regularizing assumptions, if the finite-horizon problem uses the exact state value and has the constraint that

$$x_{t+N} = 0,$$

then the MPC control law u_t^t is an asymptotically stabilizing control law. This powerful result holds for both the linear and nonlinear cases.

- In the linear case, a thorough connection between the MPC law and LQG control is established and explored in [4]. The asymptotic stability properties are relaxed from the terminal constraint to a terminal penalty. In the nonlinear case, Mayne and Michalska [10, 11] explore asymptotic stability of MPC with a reduced terminal constraint.

3. MPC FOR ABR CONGESTION CONTROL

In the ABR congestion control problem, the control delay of a single node consists of two components; the action delay in communicating to the node plus the return data travel time, and the delay in downstream data RM cells arriving at the node. Both these delays are variable. However, at any particular decision time, the delay applying to the arriving data stream is known through the arrival or non-arrival of RM cells. We may therefore pose the control problem as one in which data might be missing, but is known at the time whether it is missing or not. The state x_t of MPC is identified with target ABR ATM node queue length q_t and the allowed rate μ_t, the control u_t with command source rate $v_{i,t}$, and the measurement y_t with (q_t, μ_t).

As posed above, the MPC controller is comprised of two parts; the finite-horizon state feedback law $u_t^t = f(\hat{x}_{t|t-1})$ coming from the constrained open-loop optimal control problem, and the recursive state estimator yielding $\hat{x}_{t+1|t}$. In the linear case, such as that proposed for ABR congestion control in [1], these would be a linear state-variable feedback law and the Kalman filter — LQG control. Indeed, this is precisely the framework of [1, 3], where they adopt a certainty equivalence control approach. Here we shall explore mechanisms for accommodating delays into each of the state feedback law and the state estimator. We begin with the state estimator.

3.1. DELAYED MEASUREMENT STATE ESTIMATION

Kalman filtering theory is well developed and understood in the linear case. With a linear model, the Kalman filter may be written as the recursion of time-update and measurement update steps. The same structure pertains in nonlinear filtering using the Extended Kalman Filter [12]. These recursions have the following form.

$$\hat{x}_{t+1|t} = f(\hat{x}_{t|t}, u_t), \qquad \text{time update,} \qquad (6)$$

$$\hat{x}_{t+1|t+1} = \hat{x}_{t+1|t} + K_t(y_t - h(\hat{x}_{t+1|t})), \qquad \text{measurement update.} \qquad (7)$$

Here K_t is the Kalman gain, computed from linearizations of the system functions $f(\cdot)$, $h(\cdot)$ and the process and measurement noise covariance matrices Q and R respectively.

In the case where data y_t is unavailable, the approach is to compute successive time updates without the associated measurement update step. The same solution arises when the measurement noise covariance R is taken as infinity. In the linear case, this yields the minimum mean squared error solution. In the nonlinear case, the solution need not be optimal but still yields a workable estimator.

When the data y_t arrives with known delay, an update is still possible to generate an improved state prediction. This can be done most simply by reprocessing the data from the correctly aligned time step and then taking further time updates. Again, in the linear case this produces an MMSE optimal state estimate.

3.2. RECEDING HORIZON CONTROL WITH DELAY

The MPC control law proceeds from the current state estimate $\hat{x}_{t|t-1}$ to generate the control law

$$u_t = u_t^t(\hat{x}_{t|t-1}).$$

This controller is the first element of the receding horizon solution sequence $\{u_t^t, u_{t+1}^t, \ldots, u_{t+N-1}^t\}(\hat{x}_{t|t-1})$. If we suppose that measurements were available at time t, then in the absence of measurement data at time $t+1$ one is left with two possible options for the next control value.

Strategy 1: One could use the best currently available state estimate of $x_{t+1} = q_{t+1|t} - \mu_{t+1|t}$ and proceed to apply the stationary feedback control law u_{t+1}^{t+1}.

Strategy 2: One could apply the second element of the receding horizon solution, u_{t+1}^t.

Note that these control laws are not generically the same unless the horizon N is infinite. Both schemes do, however, rely on the same state estimate, which is produced by running the time update (6) alone without measurement update. How might one choose between these two options?

We first prove that both strategies produce asymptotically stabilizing controllers.

Theorem 1 *Suppose the receding-horizon control solution $\{u_t^t, u_{t+1}^t, \ldots, u_{t+N-1}^t\}$ is constructed by solving (4) subject to the terminal state constraint $x_{t+N} = 0$. Then each of the infinite-horizon feedback control laws $u_t = u_t^{t-k}(x_t)$ is asymptotically stabilizing for $0 \le k \le N-d$, where d is the dimension of the state x_t.*

Proof: For linear time-invariant systems one may use the monotonicity arguments of [4] (Chapter 4) with zero terminal state constraint, we see that each of the controllers may be written as $u_t = L_{N-k} x_t$ for constant LQ gain L_{N-k} derived from the Riccati difference equation (RDE) solution P_k. This RDE commences at initial condition $P_0^{-1} = 0$ and yields P_k monotonically decreasing in k for $k > d$, where d is the dimension of x_t. According to [4] Theorem 4.11, this implies the stability properties of the theorem statement. For nonlinear systems, the monotonicity of the cost-to-go serves as a Lyapunov function and the results of [9] may be used, relying on the property that $q_t = 0$ is a zero-control equilibrium. $\nabla\nabla\nabla$

To differentiate between these control strategies, we note that $u_t^{t-k}(x_t)$ is the first component in an $N-k$-step input sequence which takes x_t to the origin. Thus using u_{t+1}^{t+1} leaves us at most $N-1$ step from the origin, while using u_{t+1}^{t} leaves us at most $N-2$ steps away. Thus Strategy 2 brings us closer in time to the origin at the expense of infinite-horizon performance — provable using monotonicity arguments. Indeed, allowing the full sequence $\{u_{t+k}^{t}(x_k)\}$ to be applied takes us exactly to the origin with optimal finite-horizon performance.

But this is precisely the desired strategy for the control of ABR bit-rate in a congested network when system delays prevent output measurement — network utilization, $q_t > 0$, needs to be sacrificed for queue stability, $q_t < Q$, through the adoption of a zero-rate forcing control. The choice of horizon N may then be linked to the wind-down time required for a node to stop sending new data in response to the non-receipt of RM packets.

4. CONCLUSION

ABR congestion in high-speed networks involves coping with significant delays in asserting control action and in receiving system data. Further, we have argued that these delays will be variable when congestion is problematic. We have presented MPC as an approach to the management of feedback control in such systems through modifications to each part of its state-feedback-plus-estimator structure. For the control law, this corresponds to choosing a standard zero-terminal state constraint and to sacrificing part of the performance objective for improved stability when system measurements fail to arrive. In this way, we break with certainty equivalence design, although connections with [1] still need to be explored.

5. REFERENCES

[1] E. Altman, T. Başar, R. Srikant, "Congestion control as a stochastic control problem with action delays," *Automatica*, vol. 35, pp. 1937-1950, 1999.

[2] S. Mascolo, "Congestion control in high-speed communications netorks using the Smith principle," *Automatica*, vol. 35, pp.1921-1935.

[3] O.C. Imer, S. Compans, T. Başar, R. Srikant, "Available bit rate congestion control in ATM networks: developing explicit ratecontrol algorithms," IEEE Control Systems Magazine, vol. 21, pp. 38-56, 2001.

[4] R.R. Bitmead, M. Gevers, V. Wertz, *Adaptive Optimal Control: the thinking man's GPC*, Prentice-Hall, Sydney, 1990.

[5] E. Mosca, *Optimal, Adaptive and Predictive Control*, Prentice-Hall, Englewood Cliffs, 1995.

[6] F. Allgöwer, A. Zheng (eds), *Nonlinear Model Predictive Control*, Birkhäuser, 2001.

[7] D.Q. Mayne, J.B. Rawlings, C.V. Rao, P.O.M. Scokaert, "Constrained model predictive control: stability and optimality,"*Automatica*, vol. 36, pp. 789-814, 2000.

[8] G.J. Sutton, R.R. Bitmead, "Performance and computational implementation of nonlinear model predictive control on a submarine," in *Nonlinear Model Predictive Control*, F. Allgöwer, A. Zheng (eds), Birkhäuser, Basel, 2000.

[9] S.S. Keerthi, E.G. Gilbert, "Optimal infinite horizon feedback laws for a general class of constrained discrete time systems: stability and moving horizon approximation," *J. Optimiz Theory and Appl*, vol. 57, pp. 265-293, 1988.

[10] D.Q. Mayne and H. Michalska, "Receding horizon control of nonlinear systems," *IEEE Transaction on Automatic Control*, vol. 35, pp. 814-824, 1990.

[11] H. Michalska and D.Q. Mayne, "Robust receding horizon control of constrained nonlinear systems," *IEEE Transaction on Automatic Control*, vol. 38, pp. 1623-1632, 1993.

[12] B.D.O. Anderson, J.B. Moore, *Optimal Filtering*, Prentice-Hall US, 1979.