

Model Predictive Control for the Acquisition Queue and Related Queueing Networks

Johan S.H. van Leeuwen
EURANDOM and TU/e
Eindhoven, The Netherlands
j.s.h.v.leeuwen@tue.nl

Yoni Nazarathy
EURANDOM and TU/e
Eindhoven, The Netherlands
nazarathy@eurandom.tue.nl

Erjen Lefeber^{*}
TU/e
Eindhoven, The Netherlands
a.a.j.lefeber@tue.nl

Jacobus E. Rooda
TU/e
Eindhoven, The Netherlands
j.e.rooda@tue.nl

ABSTRACT

Model Predictive Control (MPC) is a well established method in control theory and engineering practice. It is often the method of choice for systems that need to be controlled in view of constraints. The main idea of MPC is to solve an optimization problem over a given time horizon at each control epoch, and to use the obtained solution for controlling the system until the next control epoch. In many cases this optimization problem can be formulated as a tractable quadratic programming problem.

In this paper we apply MPC to discrete-time queueing networks that incorporate multiple product routes, delays, self-generated arrivals and multi-job processing. We compare our MPC-based controller to simple threshold policies and show improved performance. We take a first exploratory step towards the application of MPC to queueing networks, highlighting the method and some open problems.

Categories and Subject Descriptors

G.3 [Probability and Statistics]: Queueing theory

General Terms

Queueing Networks, Model Predictive Control

1. INTRODUCTION

Queueing networks are often used to model communication, manufacturing and service networks. Emphasis usually lies on the performance analysis of a model for a given service discipline and routing scheme. It is often challenging to

^{*}Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

QTNA2010, July 24-26, 2010, Beijing, China.
Copyright © 2010 ACM 978-1-4503-0212-8/ ...\$10.00

determine the optimal way of scheduling capacity or routing jobs. This paper explores methods for finding such dynamic control rules.

When precise objectives are formulated, for example minimizing steady-state queue sizes, the optimal control problem can in principle be formulated as a Markov Decision Problem (MDP), but in practice, MDPs for larger networks often prove numerically cumbersome. An alternative is to employ much simpler control laws such as priority or threshold policies, but these typically fail to achieve the desired behavior. Analysis and tuning of MDPs and simple control laws can sometimes benefit from using asymptotic scaling regimes. In such cases, there is often some simplification of the underlying stochastic process, and intractable MDPs can then be approximated by more tractable Brownian control problems. Nevertheless, even when using such asymptotic approximations, finding optimal controls is still challenging. For more background see [2, 6, 11].

In this paper we present an alternative control methodology which to the best of our knowledge has not been applied before to queueing networks. Our method uses the concept of Model Predictive Control (MPC). A classic reference is [4]. MPC is a popular tool for generating feedback controllers for dynamical systems (i.e. control laws that observe the current state and use it to decide on the next control action). MPC can deal with non-linearities and state constraints. Two recent papers that apply MPC to logistical and manufacturing networks are [8] and [9], but in general, MPC is not as well known in the operations research community as it is in the systems and control community.

The concept of MPC is simple: At every control epoch solve an optimization problem for the optimal trajectory into the future, then use the first step of the optimal trajectory as the current control decision. At the next time step solve the optimization problem again, and so on. This method allows to incorporate predictions in the control mechanism and has proven useful for systems with delayed feedback. Some theoretical properties of this method have also been established [1], [10]. Specifically, conditions for obtaining a stabilizing controller in the case of deterministic systems are

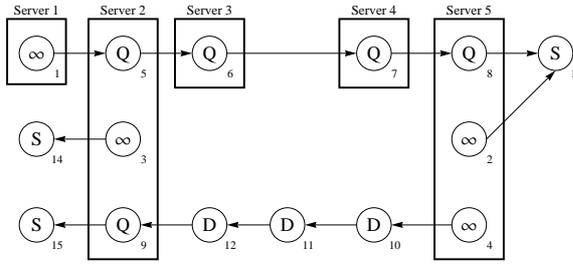


Figure 1: An example network. This network has 5 servers, 4 routes, 3 sinks, 4 infinite supply sources, 5 queues and a delay of 3 time units.

by now well known (and even implemented in commercial control packages [7]). Contrary to the deterministic case, quantifying the performance of stochastic systems controlled by MPC is a largely unexplored area.

The network models we consider are discrete-time multi-class queueing networks that allow for delays, multi-job processing and self-generated arrivals. The latter is motivated by both the acquisition queue [3] and multi-class queueing networks with infinite virtual queues [12, 15]. In these types of queueing models, a server in the network can generate arrivals and is often faced with the control choice of either letting new jobs enter the system or serving jobs that are presently waiting for service.

Figure 1 depicts an example network. The five boxes represent servers. The circles represent job classes, numbered 1 to 15. Classes marked with the ∞ sign are called *source classes*. These represent an infinite amount of jobs waiting to enter the system. Classes marked with Q represent *queue classes*; each queue class has an associated queue in which jobs are waiting to be served. Classes marked with D represent *delay classes*. When jobs pass through a delay class they are delayed for one time unit. Concatenating d delay classes results in a delay of d time units. Finally, classes marked with an S are *sink classes* and represent jobs that have left the system. We choose to model these sink classes in order to keep track of the departure processes.

Source and queue classes are associated with servers. In Figure 1, three servers are associated with a single class and two servers are associated with three classes each. All classes associated with the same server have to share that server's capacity. The delay and sink classes are not associated with servers. Jobs that finish processing in a source or queue class, or a unit delay in a delay class, move to the next *downstream* class. This is indicated by the arrows. All routes are deterministic and end up in sink classes.

The networks evolve at discrete time points $n = 0, 1, \dots$. At each time point all servers allocate their capacities among their constituent source and queue classes. The way in which this capacity is allocated is specified by the control law, the main object of our study. Capacity is used by the source classes to generate new jobs and by the queue classes to process jobs in their queue. Every unit of capacity allocated to a queue class results in the removal of a single job from

the queue. Every unit of capacity allocated to a source class results in the generation of a random number of new jobs. Delay classes simply pass, after one time unit, all of their jobs to their downstream class. Thus the randomness in our system is caused by the uncertainty in the number of jobs generated by source classes.

The structure of the paper is as follows. Section 2 illustrates the basic concepts by means of an example. Section 3 describes in detail the queueing network model, presented both as a controlled Markov chain and as a controlled linear system with noise. Section 4 describes the MPC based controller, including details of the implementation. Section 5 illustrates the applicability of MPC to arbitrary networks such as that of Figure 1. We conclude in Section 6.

An extended version of this paper is available as the technical report [13].

2. AN ILLUSTRATIVE EXAMPLE: THE ACQUISITION QUEUE

To demonstrate the MPC approach we consider the acquisition queue that was introduced in [3]. In this stochastic model, a server has to divide its capacity among the acquisition of new jobs and the service of jobs that are presently waiting in the queue. Within each time slot the server has $c \in \mathbb{N}$ units of capacity to spend. Each unit spent on acquisition in time slot n generates a random number of new jobs that join the queue after a delay of $d \in \mathbb{N}$ time units. Each unit spent on service removes one job from the queue. Figure 2 presents a schematic representation of the model.

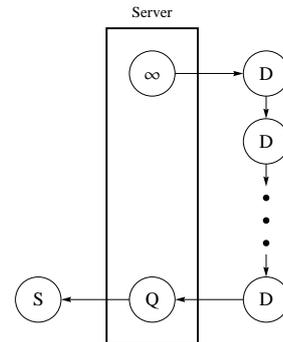


Figure 2: The acquisition queue falls within our framework. It is a single server that can either acquire new jobs (by working on the source class ∞) or service jobs in its queue. There is a delay of d time units between the time a job is acquired and its arrival to the queue. Serviced jobs accumulate in the sink.

We denote by $U(n)$ the number of acquisition efforts at time n and refer to this choice as the control law of the system. The queue length process $\{Q(n)\}_{n \in \mathbb{N}}$ is then described by the recursion equation

$$Q(n+1) = (Q(n) - (c - U(n)))^+ + \sum_{i=1}^{U(n-d)} \tilde{u}_{n,i}, \quad n = 0, 1, \dots \quad (2)$$

$$\begin{bmatrix} D_1(n+1) \\ D_2(n+1) \\ D_3(n+1) \\ Q(n+1) \\ S(n+1) \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} D_1(n) \\ D_2(n) \\ D_3(n) \\ Q(n) \\ S(n) \end{bmatrix} + \begin{bmatrix} m & 0 \\ 0 & 0 \\ 0 & 0 \\ 0 & -1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} U_\infty(n) \\ U_Q(n) \end{bmatrix} + \text{noise.} \quad (1)$$

Here, $x^+ = \max(0, x)$ and $\{\tilde{u}_{n,i}\}$ is a sequence of i.i.d. non-negative integer-valued random variables. The initial conditions are $Q(0)$ and $U(-d), U(-d+1), \dots, U(-1)$. The control $U(n) \in \mathbb{N}$ is a function of $Q(n)$ that satisfies

$$(c - Q(n))^+ \leq U(n) \leq c.$$

Let $S(n)$ be the number of jobs in the sink (jobs that have received service) by time n , with $S(0) = 0$. If we assume that the system is stable and well defined (see [3]), the throughput of the system satisfies

$$\delta = \lim_{n \rightarrow \infty} \frac{1}{n} S(n) = m \frac{c}{1+m},$$

with $m = \mathbb{E}[\tilde{u}_{1,1}]$. In [3] the following simple control law was analyzed in full detail:

$$U(n) = \alpha + (c - Q(n))^+, \quad (3)$$

with $\alpha \in \mathbb{N}$ some number for which $\alpha < c/(1+m)$ (for stability).

A sensible control law $U(\cdot)$ should stabilize the system, and in addition, keep the queue lengths relatively small. Consider the *sink error* $S^e(n) = S(n) - \delta n$. If we are able to produce at rate δ then $\lim_{n \rightarrow \infty} S^e(n)/n = 0$. The time-dependent behavior of $S^e(n)$ is also of interest: highly fluctuating or periodic behavior is considered undesirable. Finally, it is desirable for the control law to be computed relatively quickly. The threshold policy (3) certainly meets this criterion. As opposed to that, an MDP formulation of the acquisition queue (or more complex networks) typically does not. The MPC based controller presented in this paper does allow for quick evaluation.

Comparison to MPC

We now compare the threshold policy in (3) to MPC.

Figure 3(a) considers the threshold policy and shows the evolution of $Q(n)$ (above the axis) and $S^e(n)$ (below the axis). In this example, the system is deterministic (i.e. $\text{Var}(\tilde{u}_{1,1}) = 0$), and $c = 10$, $d = 10$, $m = 3$. The figure compares $\alpha = 0$ and $\alpha = 2$. It appears that $\alpha = 2$ performs better, both in terms of queue lengths and fluctuations in sink error. In Figure 3(b) we compare the threshold policy with $\alpha = 2$ and MPC¹. The latter clearly performs better. Figure 4 further compares the threshold policy and MPC, but this time for stochastic acquisition sizes. Again, MPC outperforms the threshold controller.

We now briefly outline how our MPC control for the acquisition queue comes about. For simplicity, take $d = 3$, so that the complete state at time n is described by $Q(n)$, $S(n)$, and the acquisition during the last three time units denoted by $D_1(n)$, $D_2(n)$ and $D_3(n)$.

¹With $Q = I$, $R = I$ and $N = 30$, see Section 4.

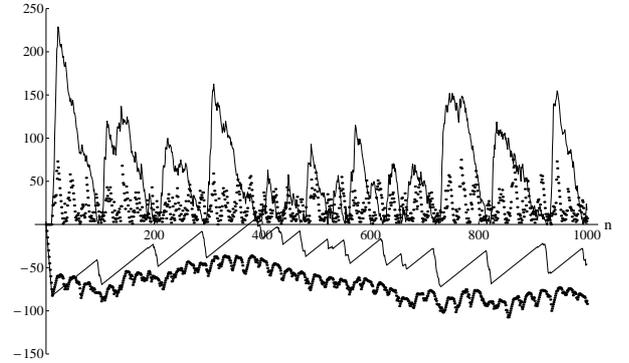


Figure 4: Trajectories of the acquisition queue with geometrically distributed acquisition sizes, $d = 10$ and $m = 3$. Trajectories above the axis are the queue level. Trajectories below the axis are the sink error. The thin solid curves are for the threshold policy with $\alpha = 2$. The dotted curves are for the MPC based controller.

Assume now that our control law is a general function of the state, and given by both the acquisition effort $U_\infty(n) \geq 0$ and the service $U_Q(n) \geq 0$, which satisfy

$$U_\infty(n) + U_Q(n) \leq c \quad \text{and} \quad U_Q(n) \leq Q(n). \quad (4)$$

The evolution of our system is as described in the linear system (1). The noise term is implicitly specified: its first coordinate is

$$\sum_{i=1}^{U_\infty(n)} \tilde{u}_{n,i} - U_\infty(n)m,$$

while all its other coordinates are zero.

The first step of our approach is to identify a *reference trajectory* by studying the system without noise. In principle we may choose any reference trajectory. For the acquisition queue, to make a fair comparison with the controller in (3), we choose as a reference the trajectory in which the sink grows linearly at rate δ and all other quantities remain constant. That is,

$$\begin{aligned} \bar{D}_1^r(n) &= \bar{D}_2^r(n) = \bar{D}_3^r(n) = \bar{Q}^r(n) = \bar{U}_\infty^r(n) = \bar{U}_Q^r(n) = \delta, \\ \bar{S}^r(n) &= \delta n. \end{aligned}$$

We now seek to formulate a *state feedback control law* that attempts to drive the system to the reference. MPC does this by solving, at each time n , the following optimization problem:

$$\begin{aligned} \min \sum_{i=n}^{n+N-1} & \left(\|\hat{X}(i+1) - \bar{X}^r(i+1)\| + \|U(i) - \bar{U}^r(i)\| \right) \\ \text{s.t.} & \quad \text{constraints (4)}. \end{aligned} \quad (5)$$

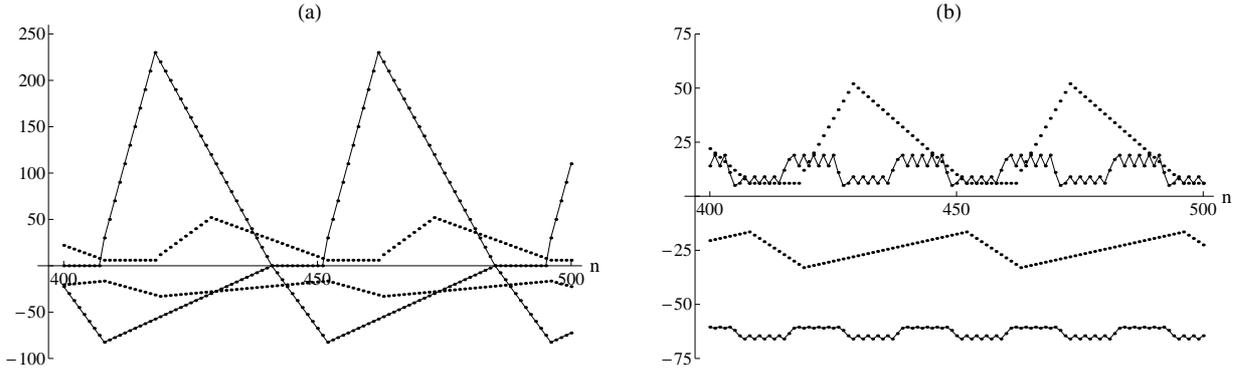


Figure 3: Trajectories of the acquisition queue with deterministic acquisition sizes, $d = 10$ and $m = 3$. Trajectories above the axis are the queue level. Trajectories below the axis are the sink error. (a) Compares the threshold policy with $\alpha = 0$ (dots with line) and $\alpha = 2$ (dots only). (b) Compares the threshold policy with $\alpha = 2$ (dots only) and an MPC based controller (dots connected by a line).

Here $\hat{X}(n)$ is the prediction of the state at time n (a vector of dimension $d + 2$ for the acquisition queue), $\bar{X}^r(\cdot)$ is the state part of the reference trajectory, $U(\cdot)$ is the vector of controls (2-dimensional for the acquisition queue) and $\|\cdot\|$ is some norm or some norm-like measure (in this paper we consider a 2-norm, i.e. quadratic costs). The minimization is over the controls which are to be applied over the time horizon $n, n + 1, \dots, n + N - 1$. This is a $2N$ -dimensional vector for the present case. The predicted state $\hat{X}(\cdot)$ is a function of the decision variables and the current state.

MPC operates as follows: At each time point the optimal solution of (5) is used to control the system for the next time unit. Note that as opposed to “myopic control” which attempts to make the next “best step”, MPC is “look ahead control”. It appears to be very suitable for queueing networks of the type we describe. The full details of the controller are surveyed in Section 4.

3. THE NETWORK MODEL

We consider multi-class queueing networks that evolve at discrete time points $n = 0, 1, 2, \dots$. A network is composed of four types of job classes: *source* (∞), *queue* (Q), *delay* (D) and *sink* (S). Jobs originate in source classes, pass through queue and delay classes, and eventually end up in sink classes.

Let K_∞, K_Q, K_D, K_S denote the number of classes of each type. Classes are indexed by $k = 1, \dots, K$. We denote the classes of type j by \mathcal{K}_j for $j \in \{\infty, Q, D, S\}$ and we number the classes as follows: $\mathcal{K}_\infty = \{1, \dots, K_\infty\}$, $\mathcal{K}_Q = \{K_\infty + 1, \dots, K_\infty + K_Q\}$, etc. For convenience, we use the notation $\mathcal{K}_{\{j, j'\}}$ to indicate the union of \mathcal{K}_j and $\mathcal{K}_{j'}$, for $j, j' \in \{\infty, Q, D, S\}$, sometimes extending the notation to three indices. For example, $\mathcal{K}_{\{\infty, Q, D\}}$ are all classes except the sink classes. Let $K_{\{j, j'\}} = |\mathcal{K}_{\{j, j'\}}|$.

The jobs are processed by L servers indexed by $i = 1, \dots, L$. These servers perform activities on the source and queue classes. The activity of each such class is performed by a

unique server $\sigma(k)$. Let

$$C(i) = \{k \in \mathcal{K}_{\{\infty, Q\}} : \sigma(k) = i\}$$

denote the constituency of server i . Let C denote the $L \times K_{\{\infty, Q\}}$ -dimensional constituency matrix. Element (i, k) of this matrix is 1 if $k \in C(i)$, and 0 otherwise. We partition C as

$$C = \begin{bmatrix} C_\infty & C_Q \end{bmatrix}.$$

At each time point, each server i must divide its effort between the constituency activities $C(i)$, generating new material from the source classes, processing existing material in queue classes, idling or performing a combination of these. The number of activities that server i can perform in one time unit is given by the integer $c_i \geq 1$. Let c denote the vector of these elements. Let $U_k(n), k \in \mathcal{K}_{\{\infty, Q\}}$ denote the number of activities that are actually applied to class k at time n . We thus have the constraints

$$\sum_{k \in C(i)} U_k(n) \leq c_i, \quad i = 1, \dots, L. \quad (7)$$

With each source class $k \in \mathcal{K}_\infty$, we associate a sequence of non-negative i.i.d. *inputs*: $\{\tilde{u}_k(\ell), \ell = 1, 2, \dots\}$, with $\mathbb{E}[\tilde{u}_k(1)] = m_k$. Let M_∞ be a diagonal matrix with the elements $m_k, k \in \mathcal{K}_\infty$. We denote generic random variables of these i.i.d. sequences by $\tilde{u}_k, k \in \mathcal{K}_\infty$. Further, use \tilde{u}_k^{*v} to denote generic random variables whose distribution is the v -fold convolution of the distribution of \tilde{u}_k . The action of performing an activity on class $k \in \mathcal{K}_\infty$ is the creation of \tilde{u}_k new jobs. Thus the application of $U_k(n)$ units of processing on class k results in the creation of $\tilde{u}_k^{*U_k(n)}$ jobs. Obviously, the mean of this quantity is $U_k(n)m_k$.

Jobs that leave a class move to downstream classes using some routing scheme. Downstream classes may be of the queue, delay, or sink type. If it is a queue class, number k say, the job joins the queue and awaits service from server $\sigma(k)$. If it is a delay class the job is delayed for one time unit, before moving onto the next downstream class in the next time unit. If it is a sink class, the job has reached its final destination.

$$X_k(n+1) = \begin{cases} X_k(n) + \sum_{k' \in \mathcal{K}_D} X_{k'}(n)p_{k'k} + \sum_{k' \in \mathcal{K}_{\{Q,\infty\}}} \tilde{u}_{k'}^{*U_{k'}(n)} p_{k'k} - U_k(n), & k \in \mathcal{K}_Q \text{ (queue),} \\ \sum_{k' \in \mathcal{K}_D} X_{k'}(n)p_{k'k} + \sum_{k' \in \mathcal{K}_{\{Q,\infty\}}} \tilde{u}_{k'}^{*U_{k'}(n)} p_{k'k}, & k \in \mathcal{K}_D \text{ (delay),} \\ X_k(n) + \sum_{k' \in \mathcal{K}_D} X_{k'}(n)p_{k'k} + \sum_{k' \in \mathcal{K}_{\{Q,\infty\}}} \tilde{u}_{k'}^{*U_{k'}(n)} p_{k'k}, & k \in \mathcal{K}_S \text{ (sink).} \end{cases} \quad (6)$$

In this paper we assume that routing is deterministic. We let $p_{kk'} = 1$ if jobs that move out of class $k \in \mathcal{K}_{\{\infty,Q,D\}}$ move into class $k' \in \mathcal{K}_{\{Q,D,S\}}$. Otherwise $p_{kk'} = 0$. These values are arranged in a matrix $P = (p_{kk'})$, which we partition as

$$P = \begin{bmatrix} P_{\infty Q} & P_{\infty D} & P_{\infty S} \\ P_{Q Q} & P_{Q D} & P_{Q S} \\ P_{D Q} & P_{D D} & P_{D S} \end{bmatrix}.$$

We assume that all routes end up in sink classes. Note that routes may merge.

The state of the network is the number of jobs in the queue, delay and sink classes. We denote it by the $K_{\{Q,D,S\}}$ -dimensional vector $X(n)$ partitioned as $[X_Q(n) X_D(n) X_S(n)]$. The state evolves according to the recursion (6). Here we use the convention $\tilde{u}_k = 1$ for $k \in \mathcal{K}_Q$, i.e. for these classes $\tilde{u}_k^{*U_k(n)} = U_k(n)$. This allows us to represent the job inflow resulting from $U_k(t)$ in the same manner for both the queue and the source classes.

Queue activities require material to be in the queue for the activity to be performed. We thus have the constraints

$$U_k(n) \leq X_k(n), \quad k \in \mathcal{K}_Q. \quad (8)$$

When the control $U(n)$ is a well specified function of the state $X(n)$ that satisfies constraints (7) and (8) we refer to it as a *state feedback control*. This makes $\{X(n)\}_{n=0}^{\infty}$ a Markov chain.

Description as a linear system with noise

As an alternative to the Markov chain representation, we now represent our network as a linear system with control-dependent zero-mean non-Gaussian noise. That is,

$$X(n+1) = AX(n) + BU(n) + G\tilde{u}(U(n)). \quad (10)$$

Here, A , B and G are matrices to be defined below and $\tilde{u}(U(n))$ denotes a K_{∞} -dimensional zero-mean random vector with the k th element distributed as

$$\tilde{u}_k^{*U_k(n)} - U_k(n)m_k.$$

Observe that the elements of this noise vector are always zero when \tilde{u}_k is deterministic. Further, when the control action $U_k(n)$ is a large number, the k th noise element is approximately Gaussian distributed.

The matrices A , B and G are spelled out in (9), I is the identity matrix.

The control $U(n)$ needs to satisfy linear constraints of the form

$$F \begin{bmatrix} X(n) \\ U(n) \end{bmatrix} \leq g,$$

with F and g further specified as

$$\begin{bmatrix} 0 & 0 & 0 & -I & 0 \\ 0 & 0 & 0 & 0 & -I \\ -I & 0 & 0 & 0 & I \\ 0 & 0 & 0 & C_{\infty} & C_Q \end{bmatrix} \begin{bmatrix} X_Q(n) \\ X_D(n) \\ X_S(n) \\ U_{\infty}(n) \\ U_Q(n) \end{bmatrix} \leq \begin{bmatrix} 0 \\ 0 \\ 0 \\ c \end{bmatrix}. \quad (11)$$

In addition we require that $U(n)$ is integer.

The fluid relaxation and reference trajectories

For the MPC approach, it is useful to consider an auxiliary system that is a deterministic fluid relaxation of (10). For this, assume that the noise component does not exist and that the control effort (and thus the queue lengths) need not satisfy integrality constraints. The resulting system is then

$$\bar{X}(n+1) = A\bar{X}(n) + B\bar{U}(n) \quad (12)$$

subject to

$$F \begin{bmatrix} \bar{X}(n) \\ \bar{U}(n) \end{bmatrix} \leq g, \quad (13)$$

with the matrices A , B , F and the vector g as defined above. For this system we define a *reference trajectory* as a pair of sequences $\{\bar{X}^r(n), \bar{U}^r(n)\}$ such that when the control $\bar{U}^r(\cdot)$ is applied to the fluid relaxation dynamics (12), the state $\bar{X}^r(\cdot)$ is obtained. In addition, the reference trajectory needs to adhere to the constraints (13). In general, the choice of the reference trajectory may depend on the desired throughput of the system. We discuss this briefly in Section 5.

Given a reference trajectory, we define the *error dynamics* as

$$X^e(n) = X(n) - \bar{X}^r(n), \quad U^e(n) = U(n) - \bar{U}^r(n). \quad (14)$$

Combining (14), (12) and (10) then yields

$$\begin{aligned} X^e(n+1) &= AX^e(n) + BU^e(n) + G\tilde{u}(U^e(n) + \bar{U}^r(n)) \\ &= AX^e(n) + BU^e(n) + \text{noise}. \end{aligned} \quad (15)$$

This representation proves useful in the next section.

4. THE MPC BASED APPROACH

We now describe the MPC approach in more detail. The controller is parameterized by a discrete *time horizon* $N > 0$ and two positive definite matrices Q and R of dimensions $K_{\{Q,D,S\}}$ and $K_{\{\infty,Q\}}$, respectively. At time n , the controller uses the optimal solution of a quadratic programming (QP) problem in which the decision variables are the controls over the time horizon: $n, n+1, \dots, n+N-1$. Given the current state $X(n)$ and the controls $U(i)$, $i = n, \dots, n+N-1$, the prediction of the state over the time horizon, denoted by $\hat{X}(\cdot)$, is generated (details below) and appears in the objective and constraints of the QP.

$$\begin{bmatrix} X_Q(n+1) \\ X_D(n+1) \\ X_S(n+1) \end{bmatrix} = \begin{bmatrix} I & P'_{DQ} & 0 \\ 0 & P'_{DD} & 0 \\ 0 & P'_{DS} & I \end{bmatrix} \begin{bmatrix} X_Q(n) \\ X_D(n) \\ X_S(n) \end{bmatrix} + \begin{bmatrix} P'_{\infty Q} M_\infty & P'_{QQ} - I \\ P'_{\infty D} M_\infty & P'_{QD} \\ P'_{\infty S} M_\infty & P'_{QS} \end{bmatrix} \begin{bmatrix} U_\infty(n) \\ U_Q(n) \end{bmatrix} + \begin{bmatrix} P'_{\infty Q} \\ P'_{\infty D} \\ P'_{\infty S} \end{bmatrix} \tilde{u}(U_\infty(n)). \quad (9)$$

With $\hat{X}^e(i) = \hat{X}(i) - \bar{X}^r(i)$ denoting the prediction of the error dynamics, we get the following QP:

$$\begin{aligned} & \min \sum_{i=n}^{n+N-1} \hat{X}^e(i+1)' Q \hat{X}^e(i+1) + U^e(i)' R U^e(i) \\ & \text{s.t.} \\ & \sum_{k \in C(j)} U_k(i) \leq c_i, i = n, \dots, N-1, j = 1, \dots, L, \\ & U_k(n) \leq X_k(n), k \in \mathcal{K}_Q, \\ & U_k(i) \leq \hat{X}_k(i), i = n+1, \dots, N-1, k \in \mathcal{K}_Q, \\ & 0 \leq U(i), i = n, \dots, N-1. \end{aligned} \quad (16)$$

The objective function in (16) quadratically penalizes deviations from the reference trajectory in both the state, using the Q matrix, and the control, using the R matrix. The first set of constraints are capacity constraints. Following that is a constraint for the control at time n not to exceed the current queue levels. It is followed by a set of constraints with respect to the predicted queue levels and with non-negativity constraints.

The QP always has a unique solution. To see this, first observe that there is always a feasible solution: $U(i) = 0, i = n, \dots, n+N-1$. Then, since Q and R are positive definite, a unique solution is guaranteed (see [14] for background on quadratic programming). Note that the solution of the QP is an $N \cdot K_{\{\infty, Q\}}$ -dimensional vector that specifies a control over the whole time horizon. From this solution, we only keep the first $K_{\{\infty, Q\}}$ coordinates to specify the control decision for the next time step. The remaining coordinates are not used. At every time point, the process is repeated: At time n , the current state is observed, the reference trajectory is calculated, and $X^e(n)$ is obtained. We denote it by X_0^e . Given a control over the time horizon, the prediction for $i > n$ is made by iterating the fluid relaxation dynamics (12). Thus $\hat{X}^e(i)$, which appears in the objective and constraints of the QP, is a function of X_0^e and the decision variables $U^e(i), i = n, \dots, n+N-1$.

We denote the first step of the optimal solution as the $K_{\{\infty, Q\}}$ -dimensional vector $U_{\text{OPT}}^e(X_0^e)$. Once $U_{\text{OPT}}^e(X_0^e)$ is determined, it is converted from the error dynamic coordinate system to the control coordinate system. We finally represent our controller as

$$U_{\text{MPC}}(n, X(n)) = \left[\min \left(\max \left([U_{\text{OPT}}^e(X(n) - X^r(n)) + U^r(n)]^{\S}, 0 \right), X_Q(n) \right) \right]_C^{\%}. \quad (17)$$

Here the max and min operators operate element-wise and $[x]^{\S}$ denotes the element-wise nearest integer to x . In addition, $[u]_C^{\%}$ denotes the following operation: Check the validity of $Cu \leq c$. Every coordinate for which the inequality is invalid represents a server in which the capacity constraint has been exceeded. In this case reduce the u by integer steps until the inequality is met, by first reducing activities for the source buffers and then reducing from the queue buffers. Within each class one can follow some arbitrary specified rule.

Note that $U_{\text{MPC}}(\cdot)$ is a function of n since the reference tra-

jectory needs to be evaluated so that the current error is plugged into $U_{\text{OPT}}^e(\cdot)$ and converted back into the coordinate system of the control. Further observe that if we take $U^r(n)$ to be constant then $U_{\text{OPT}}^e(\cdot)$ is only a function of the current error X_0^e , and does not depend on the current time.

Detailed description of the QP

We now give a detailed description of the QP. The decision variables of the QP are organized in the $N \cdot K_{\{\infty, Q\}}$ -dimensional vector \underline{U}^e . The $N \cdot K_{\{Q, D, S\}}$ -dimensional vector $\hat{\underline{X}}^e$ is the predicted state error over the time horizon $n+1, \dots, n+N$, obtained by iterating (15) and assuming no noise. The predicted error can be represented as

$$\hat{\underline{X}}^e = \underline{A} X_0^e + \underline{B} \underline{U}^e,$$

where

$$\underline{A} = \begin{bmatrix} A \\ A^2 \\ \vdots \\ A^N \end{bmatrix}, \quad \underline{B} = \begin{bmatrix} B & 0 & \dots & 0 \\ AB & B & & \vdots \\ \vdots & & \ddots & \\ A^{N-1}B & \dots & B \end{bmatrix}.$$

We need some additional matrices. Denote by \underline{Q} and \underline{R} block diagonal matrices of Q and R with dimensions $N \cdot K_{\{Q, D, S\}}$ and $N \cdot K_{\{\infty, Q\}}$, respectively. Denote by \underline{C} a block diagonal matrix of C with dimension $N \cdot L \times N \cdot K_{\{\infty, Q\}}$, and let \underline{c} be a block vector that stacks the L -dimensional c vector N times. Define

$$S_{UQ} = [0 \quad I], \quad S_{XQ} = [I \quad 0 \quad 0].$$

The $K_Q \times K_{\{\infty, Q\}}$ matrix S_{UQ} is such that when multiplied by a control vector it selects the controls for the queue classes. Similarly, S_{XQ} can be multiplied by a state vector to select the state of the queue classes. In addition, define

$$\underline{S}_{UQ}^1 = [S_{UQ} 0 \dots 0], \quad \underline{S}_{UQ}^+ = \begin{bmatrix} 0 & S_{UQ} & 0 & \dots & 0 \\ \vdots & & \ddots & & \vdots \\ \vdots & & & \ddots & 0 \\ 0 & & & & S_{UQ} \end{bmatrix},$$

$$\underline{S}_{XQ}^1 = [S_{XQ} 0 \dots 0], \quad \underline{S}_{XQ}^- = \begin{bmatrix} S_{XQ} & 0 & \dots & 0 \\ \vdots & \ddots & & \vdots \\ 0 & & S_{XQ} & 0 \end{bmatrix}.$$

Multiplying \underline{S}_{UQ}^1 by the control vector over the whole time horizon results in the control of the queue classes at the first time step. Similarly, multiplying \underline{S}_{XQ}^1 by the state vector results in the state of the queue classes at the first time step. The matrix \underline{S}_{UQ}^+ is used to select the control of the queue classes at all time steps other than the first time step. The $(N-1)K_Q \times N \cdot K_{\{Q, D, S\}}$ -dimensional matrix \underline{S}_{XQ}^- selects the queue levels at times $n+1, \dots, n+N-1$ when multiplied by the predicted state.

Given arbitrary reference trajectories, the QP is in general time-dependent. Nevertheless, the required coordinates are the controls (for both the source and queue classes) and

$$\begin{aligned}
& \min_{\underline{U}^e} \quad \underline{U}^{e'} (\underline{B}' \underline{Q} \underline{B} + \underline{R}) \underline{U}^e + 2X_0^{e'} \underline{A}' \underline{Q} \underline{B} \underline{U}^e + X_0^{e'} \underline{A}' \underline{Q} \underline{A} X_0^e \\
& \text{s.t.} \quad \begin{bmatrix} \underline{C} \\ \underline{S}_{UQ}^+ \\ \underline{S}_{UQ}^- - \underline{S}_{XQ}^- \underline{B} \\ -I \end{bmatrix} \underline{U}^e \leq \begin{bmatrix} \underline{c} \\ 0 \\ 0 \\ 0 \end{bmatrix} + \begin{bmatrix} 0 \\ \underline{S}_{XQ}^+ \\ \underline{S}_{XQ}^- \\ 0 \end{bmatrix} \underline{X}^r + \begin{bmatrix} -\underline{C} \\ -\underline{S}_{UQ}^+ \\ -\underline{S}_{UQ}^- \\ I \end{bmatrix} \underline{U}^r + \begin{bmatrix} 0 \\ \underline{S}_{XQ} \\ \underline{S}_{XQ}^- \underline{A} \\ 0 \end{bmatrix} X_0^e. \tag{18}
\end{aligned}$$

the states (for only the queue classes). In many cases it is sensible to choose constant values for these references (yielding a linear reference trajectory for the sinks), so that the QP becomes time-independent. We denote by \underline{U}^r a stacked $N \cdot K_{\{\infty, Q\}}$ -dimensional vector of the reference controls. Similarly, \underline{X}^r is a stacked $N \cdot K_{\{Q, D, S\}}$ -dimensional vector of the reference states, of which only the coordinates relevant to \mathcal{K}_Q are used in the QP.

The resulting QP is formulated in (18). It is straightforward to verify that (18) agrees with the QP (16). Formulation (18) is in the form that is required by most commercial QP solvers.

5. AN ADDITIONAL EXAMPLE

We now wish to emphasize the applicability of our method to arbitrary networks, showing how to systematically find a reference trajectory. For this we use the example of Figure 1. We employ a reference with constant queue levels and a linear increase of the sink buffers. Such a trajectory can always be found by solving a linear programming problem (LP), similar to the so-called Static Planning Problem (cf. [5]) used in multi-class queueing networks.

To find a reference trajectory, define the variables $r_i, i = 1, 2, 3, 4$ to be the long range flow rate on each of the routes. Here the route index corresponds to the source queue on which the route begins. For example, route 4 is the route passing through classes, 4, 10, 11, 12, 9 and ending in the sink buffer 15. Now use the capacity constraints (7) to indicate restrictions on the variables r_i . For example, server 5 is required to serve routes 1, 2 and 4:

$$r_1 + \frac{r_2}{m_2} + \frac{r_4}{m_4} \leq c_5.$$

Note that in the above constraint, the source classes (2 and 4) are normalized by the means m_2 and m_4 , while the queue class (class 8 on route 1) is not. Writing the constraints in this manner, together with the non-negativity constraints, defines a feasible polytope. Any point within this polytope is associated with a reference trajectory.

We find a point on the boundary of the polytope by solving an LP that minimizes $\sum_{i=1}^4 w_i r_i$. In applications, the weights w_i reflect preference of routes. This is the resulting

LP:

$$\begin{aligned}
& \max \quad \sum_{i=1}^4 w_i r_i \\
& \text{s.t.} \quad \begin{bmatrix} 1 & 0 & 0 & 0 \\ 1 & 0 & m_3^{-1} & 1 \\ 1 & m_2^{-1} & 0 & m_4^{-1} \end{bmatrix} \begin{bmatrix} r_1 \\ r_2 \\ r_3 \\ r_4 \end{bmatrix} \leq \begin{bmatrix} c_1 m_1 \wedge c_3 \wedge c_4 \\ c_2 \\ c_5 \end{bmatrix}, \\
& r_i \geq 0, \quad i = 1, 2, 3, 4. \tag{19}
\end{aligned}$$

Denote the solution by $r_i^*, i = 1, 2, 3, 4$. Further, let ρ be an arbitrary number in the interval $[0, 1]$. We now have the following reference trajectory:

$$\begin{aligned}
\bar{U}_i^r(n) &= \rho r_i^* / m_i, \quad i = 1, 2, 3, 4 \\
\bar{X}_5^r(n) &= \bar{X}_6^r(n) = \bar{X}_7^r(n) = \bar{X}_8^r(n) = \bar{U}_5^r(n) = \\
\bar{U}_6^r(n) &= \bar{U}_7^r(n) = \bar{U}_8^r(n) = \rho r_1^*, \\
\bar{X}_{10}^r(n) &= \bar{X}_{11}^r(n) = \bar{X}_{12}^r(n) = \bar{X}_9^r(n) = \bar{U}_9^r(n) = \rho r_4^*, \\
\bar{X}_{13}^r(n) &= \rho(r_1^* + r_2^*)n, \quad \bar{X}_{14}^r(n) = \rho r_3^* n, \quad \bar{X}_{15}^r(n) = \rho r_4^* n.
\end{aligned}$$

When $\rho = 0$ the reference indicates “no job flow” and when $\rho = 1$ the reference indicates “maximal job flow” (i.e. the reference is on the boundary of the polytope).

For a numerical example consider the following parameters: $c_1 = 40, c_2 = 60, c_3 = c_4 = c_5 = 30, m_{15} = m_{13} = m_{12} = 1, m_{14} = 1.8$. Use the weights: $w_1 = 10, w_2 = 5, w_3 = 1, w_4 = 10$. The solution of the LP is

$$r_1^* = 1.9081, \quad r_2^* = 38.1992, \quad r_3^* = 51.2218, \quad r_4^* = 6.8702.$$

Every $\rho \in [0, 1]$ defines a reference trajectory. We simulate with $\rho = 1$. Figure 5 plots trajectories of X_k for $k \in \mathcal{K}_Q$ and X_k^e for $k \in \mathcal{K}_S$. The parameters of the MPC based controller are $Q = I, R = I$ and $N = 3, 4, 5$. We see that when $N = 5$ both the queue levels and the sink errors are stable over a long run of 20,000 time units. For $N = 4$ we observe that X_{15}^e is not stable. Finally, when $N = 3$, both X_{15}^e and X_9 seem unstable (we only plot two coordinates in (c)). Interestingly, the trajectory of X_9 “leaves the axis” at the same time when X_{15}^e hits the axis, and from that point onwards, the two trajectories closely follow each other. This behavior appears to be consistent over several runs.

The classification of models that are able to operate at $\rho = 1$ in a stable manner is an interesting open question. To our surprise we have seen that this is possible for Example 4. We have also observed (as expected) that changing the value of the time horizon produces different results and behaviors. An alternative option is to use $N = \infty$. While in this case, the QP (18) is not well defined, there are methods for implementing such an MPC based controller, cf. [4]. The basic idea is to partition the time horizon into a finite part and an infinite part, and to use the Ricatti Equation to solve the Linear Quadratic Regulator (LQR) problem for the infinite

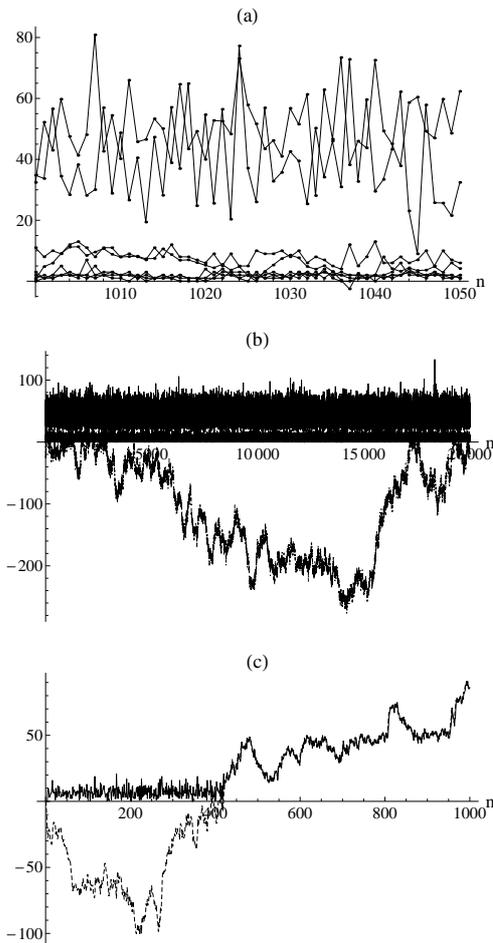


Figure 5: Trajectories of the model of Figure 1 using an MPC based controller with $Q = I$, $R = I$ and $N = 3, 4, 5$. (a) $N = 5$. In this case the system appears stable over a longer time horizon, and we only plot the trajectory over the range 1000 to 1050. (b) $N = 4$. Queue levels are stable, but the sink error of 15 is not. (c) $N = 3$. The system is not stable. Only X_{15}^e and X_9 are plotted.

part. It is possible that in this case, the structure of the controller may be found for certain Q and R matrices. This remains the subject of future work.

6. CONCLUSION

We have explored the applicability of MPC to stochastic queueing networks operating in discrete time. We have outlined a methodological way of constructing controllers for such networks. While our method appears to work well for some examples, explicit performance analysis of the behavior of the resulting stochastic processes remains an open question and is the subject of ongoing research. In this respect, one future research goal is to attempt to parallel the stability results that are currently known for deterministic settings [10] to our stochastic queueing network setting.

Of further interest is the study of queueing networks that

generate their own arrivals. A byproduct of our simulation results is the observation that the discrete time networks can often be operated on the boundary of the capacity region in a stable manner.

7. ACKNOWLEDGMENTS

This work was supported by the Netherlands Organization for Scientific Research (NWO-VIDI grant 639.072.072).

8. REFERENCES

- [1] D.P. Bertsekas. *Dynamic Programming and Optimal Control, Vol. I*. Athena Scientific, 2007.
- [2] H. Chen and D.D. Yao. *Fundamentals of Queueing Networks: Performance, Asymptotics, and Optimization*. Springer, 2001.
- [3] D. Denteneer, J.S.H. van Leeuwen, and I.J.B.F. Adan. The acquisition queue. *Queueing Systems*, 56(3):229–240, 2007.
- [4] C.E. Garcia, D.M. Preth, and M. Morari. Model predictive control: theory and practice—a survey. *Automatica*, 25(3):335–348, 1989.
- [5] J. M. Harrison. Brownian models of open processing networks: Canonical representation of workload. *Ann. Appl. Probab*, 10(1):75–103, 2000.
- [6] H.J. Kushner. *Heavy Traffic Analysis of Controlled Queueing and Communication Networks*. Springer, 2001.
- [7] M. Kvasnica, P. Grieder, M. Baotic, and M. Morari. Multi-parametric toolbox (MPT). *Hybrid Systems: Computation and Control*, pages 121–124, 2004.
- [8] M. Laumanns and E. Lefeber. Robust optimal control of material flows in demand-driven supply networks. *Physica A: Statistical Mechanics and its Applications*, 363(1):24–31, 2006.
- [9] E. Lefeber, S. Lammer, and J.E. Rooda. Optimal control of a deterministic multiclass queueing system by serving several queues simultaneously. *SE Technical Report, Systems Engineering Group, The Department of Mechanical Engineering, Eindhoven University of Technology*, 2008-09, 2008.
- [10] D.Q. Mayne, J.B. Rawlings, C.V. Rao, and P.O. Scokaert. Constrained model predictive control: stability and optimality. *Automatica*, 36:789–814, 2000.
- [11] S.P. Meyn. *Control Techniques for Complex Networks*. Cambridge University Press, 2008.
- [12] Y. Nazarathy. *On Control of Queueing Networks and the Asymptotic Variance Rate of Outputs*. PhD thesis, The University of Haifa, 2009.
- [13] J.S.H. van Leeuwen, E. Lefeber, Y. Nazarathy, and J. E. Rooda. Model predictive control for the acquisition queue and related queueing networks. *EURANDOM Technical Report Series*, 2010-22, 2010.
- [14] R.J. Vanderbei. *Linear Programming: Foundations and Extensions*. Kluwer Academic Publishers, 2001.
- [15] G. Weiss. Stability of a simple re-entrant line with infinite supply of work – the case of exponential processing times. *J. Oper. Res. Soc. Jpn.*, 47(4):304–313, 2004.