
Lecture Note 17

1 Approximate Linear Programming

Recall that we approximate $J^* \approx \Phi \tilde{r}$, where \tilde{r} is the solution of

$$\begin{aligned} \max_r \quad & c^T \Phi r \\ \text{s.t.} \quad & T \Phi r \geq \Phi r \end{aligned} \quad (\text{ALP})$$

In the previous lecture, we proved the following result on the approximation error yielded by the ALP:

Theorem 1 *If $\Phi v = e$ for some v , then we have*

$$\|J^* - \Phi \tilde{r}\|_{1,c} \leq \frac{2}{1 - \alpha} \min_r \|J^* - \Phi r\|_{\infty}. \quad (1)$$

Though the above bound offers some support for the linear programming approach, there are some significant weaknesses:

1. The bound calls for an element of the span of the basis functions to exhibit uniformly low error over all states. In practice, however, $\min_r \|J^* - \Phi r\|_{\infty}$ is typically huge, especially for large-scale problems.
2. The bound does not take into account the choice of state-relevance weights. As demonstrated in the previous section, these weights can significantly impact the approximation error. A sharp bound should take them into account.

In this lecture, we present a line of analysis that generalizes Theorem 1 and addresses the aforementioned difficulties.

2 Lyapunov Function

To set the stage for the development of an improved bound, let us establish some notation. First, we introduce a weighted maximum norm, defined by

$$\|J\|_{\infty, \gamma} = \max_{x \in \mathcal{S}} \gamma(x) |J(x)|, \quad (2)$$

for any $\gamma : \mathcal{S} \mapsto \mathfrak{R}^+$. As opposed to the maximum norm employed in Theorem 1, this norm allows for uneven weighting of errors across the state space.

We also introduce an operator H , defined by

$$(HV)(x) = \max_{a \in \mathcal{A}_x} \sum_{y \in \mathcal{S}} P_a(x, y) V(y),$$

for all $V : \mathcal{S} \mapsto \mathfrak{R}$. For any V , $(HV)(x)$ represents the maximum expected value of $V(Y)$ if the current state is x and Y is a random variable representing the next state.

For each $V : \mathcal{S} \mapsto \mathfrak{R}$, we define a scalar β_V given by

$$\beta_V = \max_x \frac{\alpha(HV)(x)}{V(x)}. \quad (3)$$

We can now introduce the notion of a ‘‘Lyapunov function,’’ as follows.

Definition 1 (Lyapunov function) *We call $V : \mathcal{S} \mapsto \mathfrak{R}^+$ a Lyapunov function if $\beta_V < 1$.*

Our definition of a Lyapunov function translates into the condition that there exist $V > 0$ and $\beta < 1$ such that

$$\alpha(HV)(x) \leq \beta V(x), \quad \forall x \in \mathcal{S}. \quad (4)$$

If α were equal to 1, this would look like a Lyapunov stability condition: the maximum expected value $(HV)(x)$ at the next time step must be less than the current value $V(x)$. In general, α is less than 1, and this introduces some slack in the condition.

Our error bound for the approximate LP will grow proportionately with $1/(1 - \beta_V)$, and we therefore want β_V to be small. Note that β_V becomes smaller as the $(HV)(x)$ ’s become small relative to the $V(x)$ ’s; β_V conveys a degree of ‘‘stability,’’ with smaller values representing stronger stability. Therefore our bound suggests that, the more stable the system is, the easier it may be for the approximate LP to generate a good scoring function.

An interesting (and useful) fact about Lyapunov functions is that T is a contraction with respect to $\|\cdot\|_{\infty, 1/V}$, when V is a Lyapunov function:

Lemma 1 *Let V be a Lyapunov function. Then, $\forall J, \bar{J}$, we have*

$$\|TJ - T\bar{J}\|_{\infty, \frac{1}{V}} \leq \beta_V \|J - \bar{J}\|_{\infty, \frac{1}{V}}$$

Proof: Let J, \bar{J} be arbitrary. Let

$$\epsilon = \|J - \bar{J}\|_{\infty, \frac{1}{V}}$$

Then

$$\bar{J} - V\epsilon \leq J \leq \bar{J} + V\epsilon$$

Let \bar{u} be greedy with respect to \bar{J} , and u be greedy with respect to J . Then

$$\begin{aligned} T\bar{J} - TJ &= (g_{\bar{u}} + \alpha P_{\bar{u}}\bar{J}) - (g_u + \alpha P_u J) \\ &\leq (g_u + \alpha P_u \bar{J}) - (g_u + \alpha P_u J) \\ &= \alpha P_u (\bar{J} - J) \\ &\leq \alpha P_u |J - \bar{J}| \end{aligned}$$

Therefore, $\forall x \in \mathcal{S}$,

$$\begin{aligned}
(TJ)(x) - (T\bar{J})(x) &= \alpha \sum_y P_u(x, y) \frac{|J(y) - \bar{J}(y)|}{V(y)} V(y) \\
&\leq \alpha \sum_y P_u(y) \underbrace{\left[\max_{y'} \frac{|J(y') - \bar{J}(y')|}{V(y')} \right]}_{\epsilon} V(y) \\
&\leq \epsilon \alpha \sum_y P_u(x, y) V(y) \\
&\leq \epsilon \alpha (HV)(x) \\
&\leq \epsilon \beta V(x)
\end{aligned}$$

Hence,

$$TJ - T\bar{J} \leq \beta_V V \epsilon$$

and we have

$$\frac{|(TJ)(x) - (T\bar{J})(x)|}{V(x)} \leq \epsilon \beta_V$$

□

We are now ready to state our main result. For any given function V mapping \mathcal{S} to positive reals, we use $1/V$ as shorthand for a function $x \mapsto 1/V(x)$.

Theorem 2 *Let \tilde{r} be a solution of the approximate LP. Then, for any $v \in \mathfrak{R}^K$ such that Φv is a Lyapunov function,*

$$\|J^* - \Phi \tilde{r}\|_{1,c} \leq \frac{2c^T \Phi v}{1 - \beta_{\Phi v}} \min_r \|J^* - \Phi r\|_{\infty, 1/\Phi v}. \quad (5)$$

Proof: Let

$$r^* = \arg \min_r \|J^* - \Phi r\|_{\infty, \frac{1}{\Phi}}$$

and

$$\epsilon = \|J^* - \Phi r^*\|_{\infty, \frac{1}{\Phi}}.$$

Let

$$\bar{r} = r^* - \epsilon \frac{1 + \beta_V}{1 - \beta_V} V$$

Then

$$\begin{aligned}
\|T\Phi \bar{r} - T\Phi r^*\|_{\infty, \frac{1}{\Phi}} &\leq \beta_V \|\Phi \bar{r} - \Phi r^*\|_{\infty, \frac{1}{\Phi}} \\
&= \beta_V \frac{\epsilon(1 + \beta_V)}{1 - \beta_V} \|\Phi V\|_{\infty, \frac{1}{\Phi}} \\
&= \beta_V \frac{\epsilon(1 + \beta_V)}{1 - \beta_V}
\end{aligned} \quad (6)$$

Moreover,

$$\|J^* - \Phi r^*\|_{\infty, \frac{1}{\Phi}} \leq \beta_V \|J^* - \Phi r^*\|_{\infty, \frac{1}{\Phi}} = \beta_V \epsilon \quad (7)$$

Thus,

$$\begin{aligned}
T\Phi\bar{r} &\geq T\Phi r^* - \epsilon\beta_V \frac{1+\beta_V}{1-\beta_V}V && \text{from (6)} \\
&\geq J^* - \beta_V\epsilon V - \beta_V\epsilon \frac{1+\beta_V}{1-\beta_V}V && \text{from (7)} \\
&\geq \Phi r^* - (1+\beta_V)\epsilon V - \beta_V\epsilon \frac{1+\beta_V}{1-\beta_V}V \\
&= \Phi\bar{r} + \epsilon \frac{1+\beta_V}{1-\beta_V}V - (1+\beta_V)\epsilon V - \beta_V \frac{1+\beta_V}{1-\beta_V}V \\
&= \Phi\bar{r}
\end{aligned}$$

Therefore, we have \bar{r} is feasible since $T\Phi\bar{r} \geq \Phi\bar{r}$. Note that

$$\|J^* - \Phi r^*\|_{\infty, \frac{1}{\beta_V}} = \epsilon \Rightarrow J^* - \epsilon V \leq \Phi r^* \leq J^* + \epsilon V.$$

We have

$$\begin{aligned}
\|J^* - \Phi\bar{r}\|_{1,c} &\leq \|J^* - \Phi\bar{r}\|_{1,c} \quad (\text{since } \tilde{r} \text{ is optimal and } \bar{r} \text{ is feasible}) \\
&= \|J^* - \Phi r^* + \frac{\epsilon(1+\beta_V)}{1-\beta_V}V\|_{1,c} \\
&\leq \|J^* - \Phi r^*\|_{1,c} + \frac{\epsilon(1+\beta_V)}{1-\beta_V}c^T V \\
&= \sum_x c(x)|J^*(x) - (\Phi r^*)(x)| + \frac{\epsilon(1+\beta_V)}{1-\beta_V}c^T V \\
&\leq \sum_x c(x)V(x)\epsilon + \frac{\epsilon(1+\beta_V)}{1-\beta_V}c^T V \\
&= \left[\epsilon + \frac{\epsilon(1+\beta_V)}{1-\beta_V}c^T V \right] c^T V \\
&= \frac{2c^T V}{1-\beta_V}\epsilon
\end{aligned}$$

□

Let us now discuss how this new theorem addresses the shortcomings of Theorem 1 listed in the previous section. We treat in turn the two items from the aforementioned list.

1. The norm $\|\cdot\|_{\infty}$ appearing in Theorem 1 is undesirable largely because it does not scale well with problem size. In particular, for large problems, the optimal value function can take on huge values over some (possibly infrequently visited) regions of the state space, and so can approximation errors in such regions.

Observe that the maximum norm of Theorem 1 has been replaced in Theorem 2 by $\|\cdot\|_{\infty, 1/\Phi v}$. Hence, the error at each state is now weighted by the reciprocal of the Lyapunov function value. This should to some extent alleviate difficulties arising in large problems. In particular, the Lyapunov function should take on large values in undesirable regions of the state space — regions where J^* is large. Hence, division by the Lyapunov function acts as a normalizing procedure that scales down errors in such regions.

2. As opposed to the bound of Theorem 1, the state-relevance weights do appear in our new bound. In particular, there is a coefficient $c^T \Phi v$ scaling the right-hand-side. In general, if the state-relevance weights are chosen appropriately, we expect that $c^T \Phi v$ will be reasonably small and independent of problem size. We defer to the next section further qualification of this statement and a discussion of approaches to choosing c in contexts posed by concrete examples.

3 Multiclass Queueing Networks

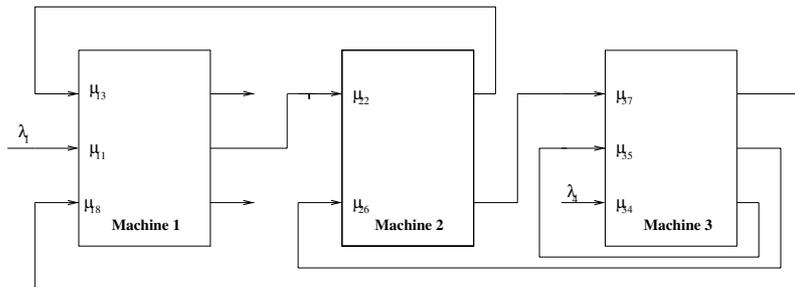


Figure 1: A Multiclass Queueing System

We now consider a queueing network with d queues and finite buffers of size B to determine the impact of dimensionality on the terms involved in the error bound of Theorem 2.

We assume that the number of exogenous arrivals occurring in any time step has expected value less than or equal to Ad , for a finite A . The state $x \in \mathfrak{R}^d$ indicates the number of jobs in each queue. The cost per stage incurred at state x is given by

$$g(x) = \frac{|x|}{d} = \frac{1}{d} \sum_{i=1}^d x_i,$$

the average number of jobs per queue.

Let us first consider the optimal value function J^* and its dependency on the number of state variables d . Our goal is to establish bounds on J^* that will offer some guidance on the choice of a Lyapunov function V that keeps the error $\min_r \|J^* - \Phi r\|_{\infty, 1/V}$ small. Since $J^* \geq 0$, we will only derive upper bounds.

Instead of carrying the buffer size B throughout calculations, we will consider the infinite buffer case. The optimal value function for the finite buffer case should be bounded above by that of the infinite buffer case, as having finite buffers corresponds to having jobs arriving at a full queue discarded at no additional cost.

We have

$$\mathbb{E}_x [|x_t|] \leq |x| + Adt,$$

since the expected total number of jobs at time t cannot exceed the total number of jobs at time 0 plus the

expected number of arrivals between 0 and t , which is less than or equal to Adt . Therefore we have

$$\begin{aligned} \mathbf{E}_x \left[\sum_{t=0}^{\infty} \alpha^t |x_t| \right] &= \sum_{t=0}^{\infty} \alpha^t \mathbf{E}_x [|x_t|] \\ &\leq \sum_{t=0}^{\infty} \alpha^t (|x| + Adt) \\ &= \frac{|x|}{1-\alpha} + \frac{Ad}{(1-\alpha)^2}. \end{aligned} \tag{8}$$

The first equality holds because $|x_t| \geq 0$ for all t ; by the monotone convergence theorem, we can interchange the expectation and the summation. We conclude from (8) that the optimal value function in the infinite buffer case should be bounded above by a linear function of the state; in particular,

$$0 \leq J^*(x) \leq \frac{\rho_1}{d}|x| + \rho_0,$$

for some positive scalars ρ_0 and ρ_1 independent of the number of queues d .

As discussed before, the optimal value function in the infinite buffer case provides an upper bound for the optimal value function in the case of finite buffers of size B . Therefore, the optimal value function in the finite buffer case should be bounded above by the same linear function regardless of the buffer size B .

As in the previous examples, we will establish bounds on the terms involved in the error bound of Theorem 2. We consider a Lyapunov function $V(x) = \frac{1}{d}|x| + C$ for some constant $C > 0$, which implies

$$\begin{aligned} \min_r \|J^* - \Phi r\|_{\infty, 1/V} &\leq \|J^*\|_{\infty, 1/V} \\ &\leq \max_{x \geq 0} \frac{\rho_1|x| + d\rho_0}{|x| + dC} \\ &\leq \rho_1 + \frac{\rho_0}{C}, \end{aligned}$$

and the bound above is independent of the number of queues in the system.

Now let us study β_V . We have

$$\begin{aligned} \alpha(HV)(x) &\leq \alpha \left(\frac{|x| + Ad}{d} + C \right) \\ &\leq V(x) \left(\alpha + \frac{\alpha A}{\frac{|x|}{d} + C} \right) \\ &\leq V(x) \left(\alpha + \frac{\alpha A}{C} \right), \end{aligned}$$

and it is clear that, for C sufficiently large and independent of d , there is a $\beta < 1$ independent of d such that $\alpha HV \leq \beta V$, and therefore $\frac{1}{1-\beta_V}$ is uniformly bounded on d .

Finally, let us consider $c^T V$. We expect that under some stability assumptions, the tail of the steady-state distribution will have an upper bound with geometric decay [1] and we take $c(x) = \left(\frac{1-\xi}{1-\xi^{B+1}} \right)^d \xi^{|x|}$. The state-relevance weights c are equivalent to the conditional joint distribution of d independent and identically

distributed geometric random variables conditioned on the event that they are all less than $B+1$. Therefore,

$$\begin{aligned}c^T V &= E \left[\frac{1}{d} \sum_{i=1}^d X_i + C \mid X_i < B+1, i = 1, \dots, d \right] \\ &< E[X_1] + C \\ &= \frac{\xi}{1-\xi} + C,\end{aligned}$$

where $X_i, i = 1, \dots, d$ are identically distributed geometric random variables with parameter $1 - \xi$. It follows that $c^T V$ is uniformly bounded over the number of queues.

References

- [1] D. Bertsimas, D. Gamarnik, and J.N. Tsitsiklis. Performance of multiclass Markovian queueing networks via piecewise linear Lyapunov functions. *Annals of Applied Probability*, 11(4):1384–1428, 2001.