

Balanced Reduced-Order Models for Iterative Nonlinear Control of Large-Scale Systems

Yizhe Huang and Boris Kramer

Abstract—We propose a new framework to design controllers for high-dimensional nonlinear systems. The control is designed through the iterative linear quadratic regulator (ILQR), an algorithm that computes control by iteratively applying the linear quadratic regulator on the local linearization of the system at each time step. Since ILQR is computationally expensive, we propose to first construct reduced-order models (ROMs) of the high-dimensional nonlinear system. We derive nonlinear ROMs via projection, where the basis is computed via balanced truncation (BT) and LQG balanced truncation (LQG-BT). Numerical experiments are performed on a semi-discretized nonlinear Burgers equation. We find that the ILQR algorithm produces good control on ROMs constructed either by BT or LQG-BT, with BT-ROM based controllers outperforming LQG-BT slightly for very low-dimensional systems.

Index Terms—Model/controller reduction, iterative learning control, large-scale systems, distributed parameter systems, fluid flow systems.

I. INTRODUCTION

HIGH-dimensional nonlinear models are ubiquitous in science and engineering, arising, e.g., during semi-discretization of partial differential equations (PDEs), or in large-scale circuits and biological models. Real-time control and filtering for these systems remains challenging. On the one hand, using linearization and subsequently linear controllers, such as the linear quadratic regulator (LQR), leads to suboptimal control for nonlinear systems. On the other hand, while nonlinear controllers can produce better control, they quickly become computationally intractable in high dimensions. The use of nonlinear reduced-order models (ROMs) is a viable alternative in these applications.

We propose a new framework to design controllers for high-dimensional nonlinear systems. The approach combines system-theoretic model reduction to obtain nonlinear ROMs and the iterative linear quadratic regulator (ILQR) control [1]. ILQR computes local linearizations at every time step, and combines it with a modified, discrete-time LQR algorithm to find the optimal controller and the corresponding nominal trajectory. The next iteration then uses the new trajectory to recompute the local linearizations and repeat the process, resulting in the iterative nature of ILQR.

Y. Huang, was with the Department of Computer Science and Engineering, University of California San Diego. He is now with the Department of Information, Risk and Operations Management, University of Texas at Austin, TX 78712 (e-mail: yizhehuang@utexas.edu)

B. Kramer is with the Department of Mechanical and Aerospace Engineering, University of California San Diego, (e-mail: bmkramer@ucsd.edu)

The ILQR algorithm has been applied to a wide range of *very low-dimensional* problems such as quadrotor UAV with cable-suspended loads [2], [3], lower limb exoskeleton models [4], and robotic arm models [1]. Moreover, several extensions to ILQR have been proposed. An iterative LQG method is developed in [5], and [6] extends ILQR to systems with non-quadratic cost functions and apply it to differential-drive robots and quadrotor helicopters in environments with obstacles. The authors in [7] propose a way of learning local linearization of dynamics from image inputs and tests it on problems like planar systems (navigation), inverted pendulum, and cart-pole balancing. In addition, [8] proposes constrained iterative LQR to handle the constraints in ILQR and applies it to on-road autonomous driving motion planning. However, while these extensions further expand the range of problems that ILQR can be applied to, none of these extensions addresses the problem that ILQR is computationally intractable for high-dimensional systems. This results from the running time of each iteration in ILQR increases quadratically with the number of dimensions.

Model reduction provides a mathematical and system-theoretic framework to reduce the dimensionality of models, see [9], [10], [11] for an overview. Here, we are interested in designing the controller using the nonlinear ROM, and using that very control on the original high-dimensional system. One of the most popular system-theoretic model reduction technique for linear open-loop systems is the balanced truncation (BT) method [12]. It finds a state-space transformation in which the controllability and observability Gramians are equal and diagonal, and then discards the states of the system that are hard to control and hard to observe. LQG balanced truncation (LQG-BT) [13] is a modification to BT; it produces ROMs that are suitable for closed-loop systems [14], [15], [16], [17]. The LQG-BT algorithm does take the matrices of the control cost function into account in that it solves the LQG Riccati equations instead of the Lyapunov equations. In this work, we compare the performance of ILQR controllers based on nonlinear ROMs, which we compute via BT and LQG-BT model reduction. Our numerical testbed is the semi-discretized high-dimensional Burgers equation.

This paper is organized as follows. Section II briefly introduces the ILQR algorithm, presents algorithmic complexity results, and the special case of quadratic-in-state systems. Section III presents BT and LQG-BT model reduction. Section IV numerically investigates the performance of the BT and LQG-based controllers. Section V offers conclusions and an outlook for future work.

II. ITERATIVE LINEAR QUADRATIC REGULATOR

We define the nonlinear control problem in Section II-A. Then, in Section II-B we introduce the ILQR method, first developed in [1], which extends LQR to nonlinear systems. In Section II-C we derive a computational cost analysis of ILQR, and in Section (II-D) we specifically analyze quadratic-in-state systems in combination with ILQR.

A. Nonlinear Control Problem

We consider the nonlinear control problem of minimizing a quadratic cost subject to a nonlinear dynamical system:

$$\begin{aligned} \min_{u_0, u_1, \dots, u_{N-1}} \quad & J(x_{0:N}, u_{0:N-1}) \\ \text{s.t.} \quad & x_{k+1} = f(x_k, u_k). \end{aligned} \quad (1)$$

where $x_k \in \mathbb{R}^n$ is the system state at time step k , $u_k \in \mathbb{R}^m$ is the control input at time step k , N is the number of time steps, $x_{0:N} = \{x_0, x_1, \dots, x_N\}$ and $u_{0:N-1} = \{u_0, u_1, \dots, u_{N-1}\}$ denote the sequence of states and controls, respectively, and $f: \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n$ is a differentiable function in both arguments that maps the current state and control to the future state. The cost function is assumed in quadratic form

$$J(x_{0:N}, u_{0:N-1}) = (x_N - x^*)^\top Q_f (x_N - x^*) + \sum_{k=0}^{N-1} (x_k^\top Q x_k + u_k^\top R u_k), \quad (2)$$

where $x^* \in \mathbb{R}^n$ is the target state. The symmetric positive semi-definite matrix $Q_f \in \mathbb{R}^{n \times n}$ defines the final state cost, and the symmetric positive semi-definite matrix $Q \in \mathbb{R}^{n \times n}$ defines the intermediate state cost, while the positive definite matrix $R \in \mathbb{R}^{m \times m}$ defines the control cost.

B. The ILQR Algorithm

The ILQR algorithm provides a solution for the nonlinear control problem in (1). The iterative algorithm starts with a nominal control sequence $u_{0:N-1}$ and nominal trajectory $x_{0:N}$ that results from that input. The algorithm then improves $u_{0:N-1}$ at each iteration so to minimize (2) while using local linearizations at every time step. We refer to [1] for details. ILQR iteratively computes a local linearization

$$\delta x_{k+1} = A_k \delta x_k + B_k \delta u_k \quad (3)$$

given a nominal control sequence $u_{0:N-1}$ and corresponding trajectory $x_{0:N}$, where δx_k and δu_k denote the deviation from the nominal trajectory, while $A_k = D_x f(x_k, u_k)$ and $B_k = D_u f(x_k, u_k)$ are the Jacobians of f with respect to x and u , respectively. The algorithm then uses the local linearization in a modified LQR algorithm to compute a local optimal control, which is

$$\delta u_k = -K_k \delta x_k - K_k^v v_{k+1} - K_k^u u_k, \quad (4)$$

where $K_k^v = (B_k^\top S_{k+1} B_k + R)^{-1} B_k^\top$ and $K_k^u = (B_k^\top S_{k+1} B_k + R)^{-1} R$, and where

$$\begin{aligned} K_k &= (B_k^\top S_{k+1} B_k + R)^{-1} B_k^\top S_{k+1} A_k \\ S_k &= A_k^\top S_{k+1} (A_k - B_k K_k) + Q \\ v_k &= (A_k - B_k K_k)^\top v_{k+1} - K_k^\top R u_k + Q x_k. \end{aligned}$$

Observe that to compute S_k and v_k we need S_{k+1} and v_{k+1} , yielding a backward-in-time iteration. In addition, we have that $S_N = Q_f$, $v_N = Q_f (x_N - x^*)$. This control is then used to compute the nominal trajectory in the next iteration, and the process proceeds until convergence. Algorithm 1 summarizes the ILQR algorithm. We terminate the algorithm when the relative difference of cost between two consecutive iterations is small, $|J_{\text{old}} - J|/J_{\text{old}} \leq \text{tol}$.

Algorithm 1: ILQR Algorithm

```

1 Input: Nonlinear term  $f(x_k, u_k)$ , initial state  $x_0$ , target
    $x^*$ , initial control sequence  $u_{0:N-1}$ , number of time
   steps  $N$ , matrices in cost function  $Q$ ,  $Q_f$ ,  $R$ ,
   convergence threshold  $\text{tol}$ ;
2 Output: Control  $u_{0:N-1}$ , number of ILQR iterations  $l$ ;
3  $l = 0$ ,  $S_N = Q_f$ ,  $\delta x_0 = 0$ ;
4  $x_{k+1} = f(x_k, u_k)$ ,  $\forall k = 0, 1, \dots, N-1$ ;
5  $J_{\text{old}} = 1$ ,  $J = 1 + 2\text{tol}$ ;
6 while  $|J_{\text{old}} - J|/J_{\text{old}} > \text{tol}$  do
7    $J_{\text{old}} = J$ ;
8    $J = 0$ ;
9    $v_N = Q_f (x_N - x^*)$ ;
10   $x_{1:N}^{\text{old}} = x_{1:N}$ ;
11  for  $k \leftarrow 0$  to  $N-1$  do
12     $A_k = D_x f(x_k, u_k)$ ,  $B_k = D_u f(x_k, u_k)$ 
13  end
14  for  $k \leftarrow N-1$  to  $0$  do
15     $K_k = (B_k^\top S_{k+1} B_k + R)^{-1} B_k^\top S_{k+1} A_k$ ;
16     $K_k^v = (B_k^\top S_{k+1} B_k + R)^{-1} B_k^\top$ ;
17     $K_k^u = (B_k^\top S_{k+1} B_k + R)^{-1} R$ ;
18     $S_k = A_k^\top S_{k+1} (A_k - B_k K_k) + Q$ ;
19     $v_k = (A_k - B_k K_k)^\top v_{k+1} - K_k^\top R u_k + Q x_k$ ;
20  end
21  for  $k \leftarrow 0$  to  $N-1$  do
22     $\delta u_k = -K_k \delta x_k - K_k^v v_{k+1} - K_k^u u_k$ ;
23     $u_{k+1} = u_k + \delta u_k$ ;
24     $x_{k+1} = f(x_k, u_k)$ ;
25     $\delta x_{k+1} = x_{k+1}^{\text{old}} - x_{k+1}$ ;
26     $J = J + x_{k+1}^\top Q x_{k+1} + u_{k+1}^\top R u_{k+1}$ ;
27  end
28   $J = J + x_N^\top Q_f x_N$ ,  $l = l + 1$ ;
29 end

```

C. Computational Cost of ILQR

In each iteration, Algorithm 1 completes three portions. First, we obtain the linearizations (step 11-13); second, we compute $K_k, K_k^v, K_k^u, S_k, v_k$ (step 14-20); third, we compute the control increment $(\delta u)_{0:N-1}$ and the new trajectory $x_{0:N}$ (step 21-27). Linearizations are obtained in N Jacobian calls if it is known analytically (e.g., for polynomial models). If we need to approximate the Jacobian numerically, then we have to march the system equation once for each dimension of state and control, which requires evaluating $f(x_k, u_k)$ at total of $N(n+m)$ times. The second and third portion require matrix multiplication and inversion. State-of-the-art algorithms

of those operations have complexity $\mathcal{O}(n^{2.4})$ [18], whereas classical methods scale as $\mathcal{O}(n^3)$, so the complexity of the second and third pass is at best $\mathcal{O}(N(n^{2.4} + m^{2.4}))$.

Proposition 1: The complexity of the l th iteration of the ILQR Algorithm 1 is $\mathcal{O}(N(n^{2.4} + m^{2.4}))$, where n is the state dimension, m is the number of controls, and N is the number of time steps.

The previous result analyzes a single iteration of the ILQR algorithm. We note that the number of iterations of the ILQR algorithm l varies significantly, depending on the structure of the nonlinear function $f(x_k, u_k)$, as we see next.

Proposition 2: When applied to a discrete-time (DT) LTI system, the ILQR algorithm converges in one iteration. The resulting control is the same as the output of DT-LQR¹.

D. ILQR for Quadratic-in-State Systems

The Euler equations in specific volume form describe a purely quadratic system and are commonly used in aerospace applications. Other models, such as Navier-Stokes and Burgers equations are nearly quadratic for small viscous terms. Thus, we analyze the ILQR algorithm specifically for quadratic-in-state systems of the form:

$$x_{k+1} = G(x_k \otimes x_k) + Bu_k, \quad (5)$$

where $G \in \mathbb{R}^{n \times n^2}$, $B \in \mathbb{R}^{n \times m}$ and \otimes denotes the Kronecker product, defined as $x \otimes x = [x_k^{(1)} x_k^\top, \dots, x_k^{(n)} x_k^\top]^\top \in \mathbb{R}^{n^2}$.

Proposition 3: Let $Q = 0$ in the cost equation (2) (penalty only on the final state). Suppose the initial state is $x_0 = 0$, and the initial control sequence is $u_{0:N-1}$. Then the ILQR algorithm applied to the quadratic system (5) takes at least N iterations to converge, where N is the number of time steps.

Proof: Since $x_0 = 0$ and initial control $u_k = 0, k = 1, 2, \dots$, the initial trajectory of the system (5) is stationary: $x_k = 0$ for $k = 0, 1, \dots, N-1$. Likewise, in the local linearization $A_k = D_x(G(x_k \otimes x_k) + Bu_k) = G(x_k \otimes I_n + I_n \otimes x_k)$, therefore, $A_k = 0$ in all local linearizations and $B_k = D_u(G(x_k \otimes x_k) + Bu_k) = B$. We first analyze the behavior of the ILQR algorithm in the first iteration. There, except S_N , all of $S_k = 0$, since $A_k = Q = 0$. Therefore, with the exception of step $N-1$, we have $K = 0$, $K_v = R^{-1}B^\top$ and $K_u = I$. Therefore $(A_k - B_k K) = 0 \forall k$, which means that all $v_k = 0$ except v_N . Therefore, $\delta u_k = 0$ except δu_{N-1} . In particular, $\delta u_{N-1} = (B^\top Q_f B + R)^{-1} B^\top Q_f x^*$. As a result, only u_{N-1} is updated in the first iteration. Similarly, we find that only u_{N-2} and u_{N-1} are updated in the second iteration, and so on. It takes N iterations before u_0 gets updated; thus, it takes at least N iterations for ILQR algorithm to converge. ■

The previous theorem points to a computational challenge. If it takes $O(N)$ iterations to converge and since in each iteration the time complexity is $O(N)$ with respect to the number of time steps (see Proposition 1) the total time complexity is $O(N^2)$. For stiff problems and long simulation times this becomes intractable. To address this problem, it may be helpful to use a coarse time step on the first few iterations and then use finer timer step later to speed up convergence. For example,

say $N = 10$ for the first few iterations, then all u_k can be updated in 10 iterations, and we can later switch to $N = 1000$. This strategy can be thought of as using the output from coarse ILQR to initialize fine-grain ILQR. Clearly, the use of ROMs also speeds up the algorithm since the time needed for each iteration is reduced significantly.

III. SYSTEM-THEORETIC MODEL REDUCTION

We present two system-theoretic model reduction techniques: balanced truncation [12] in Section III-A, and LQG-BT [13] in Section III-B. Both methods are defined for continuous-time LTI systems $\dot{x} = Ax + Bu$, $y = Cx$ where $y \in \mathbb{R}^p$ is the output and the system matrices are $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times m}$ and $C \in \mathbb{R}^{p \times n}$. We use these methods for discrete-time nonlinear systems as follows. First, we linearize the nonlinear systems and then convert the system from discrete-time to continuous-time using a bilinear Tustin transformation. We then use BT and LQG-BT to compute the state-space transformations, and use those transformations to reduce the nonlinear systems. The control inputs keep the system around the equilibrium/linearization point. Thus, using these methods to obtain the projection matrices for the nonlinear ROM is more appropriate than in the open-loop nonlinear setting. More ROM details are given in Section IV-B. Both concepts are for infinite-time horizon. There are finite-time BT methods [20], yet we found the infinite-time BT method to outperform the finite-time BT on our test problems.

A. Computing the Balanced Truncation Basis

Balanced truncation computes a state-space transformation for stable systems in such a way that the controllability and observability Gramians of the ROM are equal and diagonal, e.g., only the easy-to-control and easy-to-observe states are kept in the ROM. The controllability Gramian P , and the observability Gramian Q , satisfy the Lyapunov equations

$$AP + PA^\top + BB^\top = 0, \quad QA + A^\top Q + C^\top C = 0.$$

Applying the transformation $x = T\tilde{x}$ to the LTI system yields $\dot{\tilde{x}} = \tilde{A}\tilde{x} + \tilde{B}u$, $\tilde{y} = \tilde{C}\tilde{x}$ where $\tilde{A} = T^{-1}AT$, $\tilde{B} = T^{-1}B$, $\tilde{C} = CT$. The Gramians of the transformed system are $\tilde{P} = T^{-1}PT^{-\top}$ and $\tilde{Q} = T^\top QT$. We want T such that $\tilde{P} = \tilde{Q} = \Sigma$, hence

$$\Sigma^2 = \tilde{P}\tilde{Q} = T^{-1}PT^{-\top}T^\top QT = T^{-1}PQT.$$

Therefore, T can be obtained by computing the eigenvectors of PQ and eigenvalues σ_i^2 in $\Sigma^2 = \text{diag}(\sigma_1^2, \dots, \sigma_n^2)$. To ensure $\tilde{P} = \tilde{Q}$, proper scaling of the eigenvectors is required. The complete balanced truncation method for stable systems in an efficient implementation is given in Algorithm 2. This implementation only computes the leading r columns of T_r , and T_l and is therefore numerically more stable and efficient.

B. Computing the LQG-Balanced Truncation Basis

LQG-BT [13] applies BT to closed-loop (stabilized) systems, and therefore also applies to unstable systems where, however, (A, B) is stabilizable and (A, C) is detectable. It

¹Refer to, e.g., [19, Thm. 6.28] for DT-LQR control.

Algorithm 2: Balanced truncation algorithm

Input: Matrices $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times m}$ and $C \in \mathbb{R}^{p \times n}$,
ROM dimension r

Output: Transformation matrices T_r, T_l

- 1: Solve $AP + PA^\top + BB^\top = 0$ for $P = RR^\top$
 - 2: Solve $QA + A^\top Q + C^\top C = 0$ for $Q = LL^\top$
 - 3: $r = \min(r, \text{rank}(L), \text{rank}(R))$
 - 4: $U, \Sigma, V = \text{svd}(L[:, 1:r]^\top R[:, 1:r])$
 - 5: $T_r = R[:, 1:r]V\Sigma^{-\frac{1}{2}}$
 - 6: $T_l = \Sigma^{-\frac{1}{2}}U^\top L[:, 1:r]^\top$
-

balances the unique positive definite solutions P, Q of the LQG algebraic Riccati equations

$$\begin{aligned} AP + PA^\top - PC^\top CP + BB^\top &= 0, \\ A^\top Q + QA - QBR^{-1}B^\top Q + C^\top C &= 0. \end{aligned}$$

LQG-BT follows the same steps as Algorithm 2, except at Steps 1–2 it uses the solutions to the LQG AREs. As a byproduct, we directly obtain LQR or LQG controllers.

IV. NUMERICAL RESULTS: BURGERS EQUATION

Burgers equation has been widely used in modeling of fluids and traffic flows, and as a numerical testbed for nonlinear model reduction and control [21], [22], [23], [24].

A. PDE Model and Semi-Discretized Model

We consider a 1D Burgers equation with the setup as in [25]. The PDE model is

$$\dot{z}(\xi, t) = \varepsilon z_{\xi\xi}(\xi, t) - \frac{1}{2}(z^2(\xi, t))_{\xi} + \sum_{k=1}^m \chi_{[(k-1)/m, k/m]}(\xi) u_k(t)$$

for $t > 0$ and where ξ is the spatial variable, $u_k(t)$ is the distributed control, and m is the dimension of control, $z \in H_{\text{per}}^1(0, 1)$, which means that the system has periodic boundary conditions. Moreover, $z(\xi, 0) = z_0(\xi) = 0.5 \sin(2\pi\xi)^2$ for $\xi \in [0, 0.5]$ and zero otherwise. The function $\chi_{[a,b]}(x)$ is the characteristic function over $[a, b]$. Thus, the spatial domain is divided into m equidistant intervals, and each control is applied on one corresponding interval. When discretized with linear finite elements, and after an inversion with the mass matrix, the semi-discretized system is a nonlinear ODE which is quadratic in the state, and linear in the control:

$$\dot{x} = Ax + G(x \otimes x) + Bu, \quad y = Cx, \quad (6)$$

where $G \in \mathbb{R}^{n \times n^2}$, and initial condition $x(0) = x_0$. We choose the following parameters for our numerical experiments: $n = 101$ states, $m = 5$ controls, and viscosity $\varepsilon = 5 \times 10^{-3}$ to make the nonlinear quadratic term dominant. The system is simulated for $t_f = 5$ s, and the time interval is divided into $N = 500$ time steps. The trajectory is computed using the backward Euler method, which is implemented using the Levenberg-Marquardt algorithm. The output matrix $C = I_n$, the state cost matrices $Q = Q_f = C^\top C$, and the input cost matrix is $R = 10^3 \times I_m$, where I_n is the $n \times n$ identity matrix.

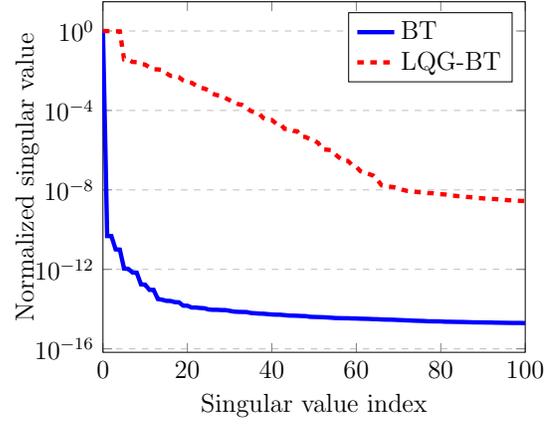


Fig. 1: Normalized singular values of the matrix $L^\top R$. For balanced truncation, those are the Hankel singular values.

B. BT and LQG-BT Reduced-Order Models

We linearize (6) around the zero equilibrium and use the linearized matrices A_{lin}, B, C in the BT or LQG-BT Algorithm 2 to compute the transformations T_r, T_l . The linearized matrix A_{lin} is stable. We then project the full-order model (FOM) in (6) to obtain

$$\dot{x}_r = A_r x_r + G_r(x_r \otimes x_r) + B_r u_r, \quad y_r = C_r x_r \quad (7)$$

where $A_r = T_l A T_r$, $B_r = T_l B$, $C_r = C T_r$, $G_r = T_l G(T_r \otimes T_r)$ and $x_r(0) = T_l x_0$. Fig. 1 shows the normalized singular values of the matrix $L^\top R$ in the balancing procedure (see step 4 in Algorithm 2) for both BT and LQG-BT. For the standard BT algorithm, those are also known as the Hankel singular values. We observe that the first 20 Hankel singular values decay fast, and all Hankel singular values except the first one come in pairs. In contrast, the normalized singular values in the LQG-BT decay much more slowly, and only level off at $r > 60$. Similar to the Hankel singular values, these singular values except the first one come in pairs. The fact that the singular values in the BT algorithm decay much faster indicates that it is likely that BT outperforms LQG-BT in open-loop simulation accuracy, as when we truncate the transform matrix we are discarding dimensions that have lower associated singular values.

C. Performance of ROM-based ILQR Controllers

We use the ROMs to compute controllers u_r^{BT} and u_r^{LQGBT} via the ILQR Algorithm 1 with $\text{tol} = 3 \times 10^{-5}$. We note that the Jacobians of (7), which are used to obtain the local linearization of the system in the ILQR algorithm, are $D_{x_r} \dot{x}_r = A_r + G_r(x_r \otimes I_{n_r} + I_{n_r} \otimes x_r)$ and $D_{u_r} \dot{x}_r = B_r$, where n_r is the number of dimensions of the state x_r in the ROM. The availability of the Jacobians in explicit form significantly speeds up the computational routine, and should be used whenever available, c.f. Proposition 1.

Fig. 2 shows the ROM cost $J(x_r, u_r)$ computed as in (2) yet with state x_r , controls u_r and cost matrix $Q_r = Q_{f,r} = C_r^\top C_r$, versus the iterations of the ILQR algorithm. Moreover, the figure shows the FOM cost $J(x, u_r)$ from (2) when the ROM

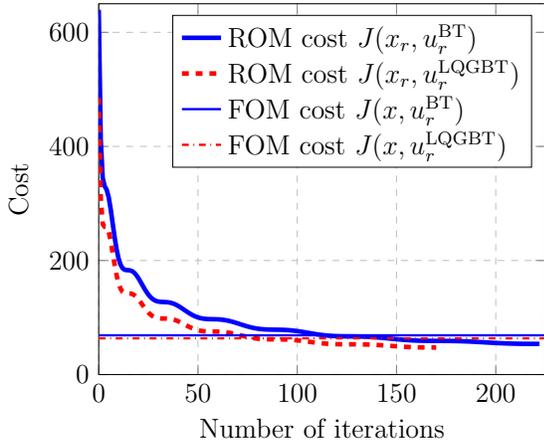


Fig. 2: ROM cost vs number of iterations in the ILQR algorithm for $r = 5$. The FOM cost when using the converged ILQR controller shown in thin lines, which is 68.9 for BT and 63.6 for LQGBT.

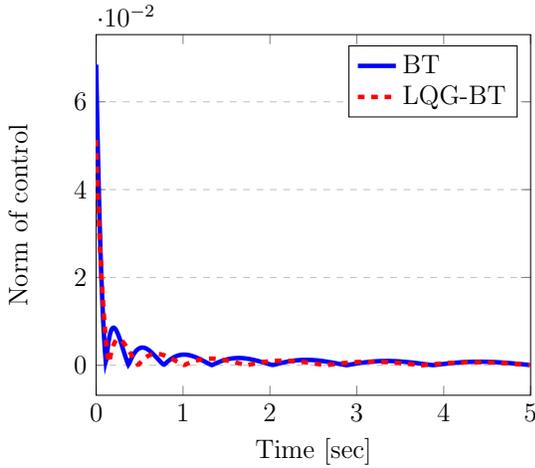


Fig. 3: Norm of ILQR control $\|u\|_2$ computed using the two nonlinear ROMs with $r = 5$.

controller is used for the FOM (thin horizontal lines). Both ROMs had dimension $r = 5$. Note that the ILQR algorithm converges monotonically for both models terminating requiring 170 iterations with the LQG-BT model and 223 iterations with the BT model. With the final controllers the cost function evaluated for the FOM, $J(x, u_r)$, is 68.9 for BT and 63.6 for LQG-BT showing a slightly better performance of LQG-BT. The norm of the controls obtained from ILQR is shown in Fig. 3, and Fig. 4 shows the corresponding output. Due to the periodic boundary conditions and the large penalty on the control, the state shows a low-magnitude periodic behavior.

Fig. 5 shows the open-loop behavior of the system when excited with the initial condition x_0 and no inputs. The initial condition gets convected in a nonlinear fashion to the right. In contrast, the controlled systems with the ROM-based ILQR controllers are shown in Fig. 6a for BT and Fig. 6b for LQG-BT. In both cases, the control signal (Fig. 3) reduces the initial condition significantly, showing the efficacy of the control.

Finally, while the above results were obtained for $r = 5$,

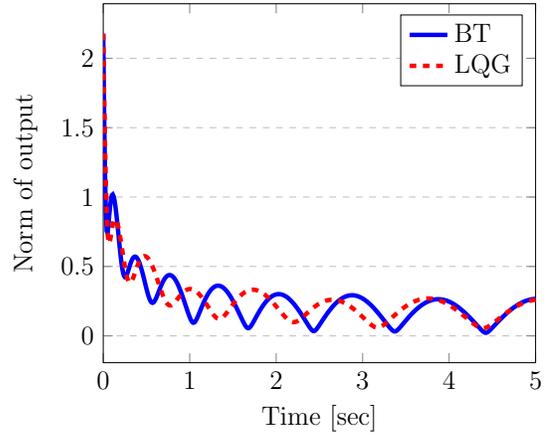


Fig. 4: Norm of outputs $\|y\|_2$ in the closed-loop FOM, with the ILQR controller computed using the two nonlinear ROMs with $r = 5$.

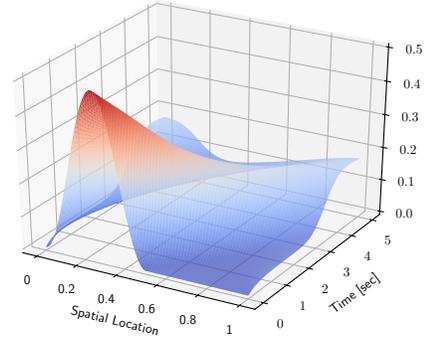
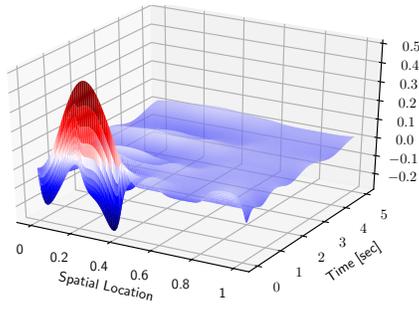


Fig. 5: FOM simulation: initial condition x_0 and zero input.

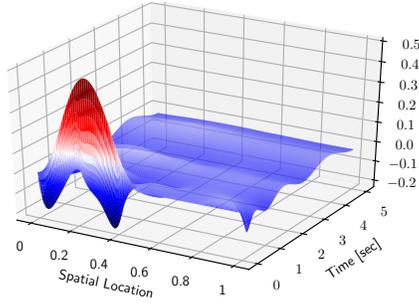
Table I gives an overview of results for different dimensions r of the nonlinear ROMs. We note that while there is an anomaly at $r = 9$ for the BT-ROM-based controller, in general, BT-ROM controllers seem to outperform the LQG-BT ROM controllers, evidenced by a lower number of iterations, as well as FOM cost, which appears to be a result of the slow singular value decay of the LQG-BT singular values in Fig. 1.

V. CONCLUSIONS AND FUTURE DIRECTIONS

We presented a new control framework for high-dimensional nonlinear systems by combining the iterative linear quadratic regulator (ILQR) framework and system-theoretic model reduction. As shown in the numerical experiments, both standard balanced truncation model reduction and LQG balanced truncation allowed us to compute well-performing controllers from ROMs with about 95% reduction of model order compared to the full-order model. Further work may be done by using nonlinear system-theoretic model reduction, such as [26], [27], [28], which despite being open-loop techniques, can provide more accurate ROMs for nonlinear systems. One would also need to investigate the observed correlation between larger state dimensions and increased iteration numbers for ILQR to understand scalability of ILQR. Lastly, convergence of ILQR for increasing N remains an open problem.



(a) Standard BT



(b) LQG-BT

Fig. 6: FOM simulation: initial condition x_0 , controlled by u_r^{BT} and u_r^{LQGBT} with $r = 5$.

TABLE I: Comparison of cost of ROM controllers on FOM, $J(x, u_r)$, on ROM, $J(x_r, u_r)$, and number of iterations of ILQR for increasing dimension of the ROM.

r	LQG-BT $J(x_r, u_r)$	LQG-BT $J(x, u_r)$	BT $J(x_r, u_r)$	BT $J(x, u_r)$	LQG-BT iter.	BT iter.
2	56.8	197.3	85.9	245.0	89	58
3	71.2	136.6	44.5	101.8	58	278
4	45.8	89.8	46.6	83.0	170	277
5	47.4	63.6	53.8	68.9	170	222
6	50.5	69.0	44.3	47.7	1328	966
7	50.8	61.3	43.9	46.5	1208	861
8	54.2	63.3	43.5	45.7	3647	930
9	53.3	59.1	272.7	272.1	9070	101
10	53.8	58.8	51.5	51.2	2395	2687
11	60.4	60.2	48.8	48.7	7927	4165

ACKNOWLEDGMENT

B.K. thanks Prof. Hermann Mena for discussions on ILQR.

REFERENCES

- [1] W. Li and E. Todorov, "Iterative linear quadratic regulator design for nonlinear biological movement systems." in *International Conference on Informatics in Control, Automation and Robotics*, 2004, pp. 222–229.
- [2] Y. Alothman and D. Gu, "Quadrotor transporting cable-suspended load using iterative linear quadratic regulator (ILQR) optimal control," in *2016 8th Computer Science and Electronic Engineering (CEECE)*. IEEE, 2016, pp. 168–173.
- [3] Y. Alothman, M. Guo, and D. Gu, "Using iterative LQR to control two quadrotors transporting a cable-suspended load," *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 4324–4329, 2017.
- [4] J. Sergey, S. Sergei, and Y. Andrey, "Comparative analysis of iterative LQR and adaptive PD controllers for a lower limb exoskeleton," in *2016 IEEE International Conference on Cyber Technology in Automation, Control, and Intelligent Systems (CYBER)*. IEEE, 2016, pp. 239–244.

- [5] E. Todorov and W. Li, "A generalized iterative LQG method for locally-optimal feedback control of constrained nonlinear stochastic systems," in *Proceedings of the 2005, American Control Conference, 2005*. IEEE, 2005, pp. 300–306.
- [6] J. van den Berg, "Iterated LQR smoothing for locally-optimal feedback control of systems with non-linear dynamics and non-quadratic cost," in *2014 American Control Conference*. IEEE, 2014, pp. 1912–1918.
- [7] E. Banijamali, R. Shu, H. Bui, A. Ghodsi *et al.*, "Robust locally-linear controllable embedding," in *International Conference on Artificial Intelligence and Statistics*, 2018, pp. 1751–1759.
- [8] J. Chen, W. Zhan, and M. Tomizuka, "Constrained iterative LQR for on-road autonomous driving motion planning," in *2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2017, pp. 1–7.
- [9] A. C. Antoulas, *Approximation of Large-Scale Dynamical Systems*. SIAM, 2005.
- [10] A. Quarteroni and G. Rozza, *Reduced order methods for modeling and computational reduction*. Springer, 2014.
- [11] P. Benner, M. Ohlberger, A. Cohen, and K. Willcox, *Model Reduction and Approximation*. Society for Industrial and Applied Mathematics, 2017.
- [12] B. Moore, "Principal component analysis in linear systems: Controllability, observability, and model reduction," *IEEE Transactions on Automatic Control*, vol. 26, no. 1, pp. 17–32, 1981.
- [13] E. Jonckheere and L. Silverman, "A new set of invariants for linear systems—application to reduced order compensator design," *IEEE Transactions on Automatic Control*, vol. 28, no. 10, pp. 953–964, 1983.
- [14] B. B. King, N. Hovakimyan, K. A. Evans, and M. Buhl, "Reduced order controllers for distributed parameter systems: LQG balanced truncation and an adaptive approach," *Mathematical and computer modelling*, vol. 43, no. 9-10, pp. 1136–1149, 2006.
- [15] J. R. Singler and B. A. Batten, "A comparison of balanced truncation methods for closed loop systems," in *2009 American Control Conference*. IEEE, 2009, pp. 820–825.
- [16] T. Breiten and K. Kunisch, "Feedback stabilization of the Schlögl model by LQG-balanced truncation," in *2015 European Control Conference (ECC)*. IEEE, 2015, pp. 1171–1176.
- [17] P. Benner and J. Heiland, "LQG-balanced truncation low-order controller for stabilization of laminar flows," in *Active Flow and Combustion Control 2014*. Springer, 2015, pp. 365–379.
- [18] D. Coppersmith and S. Winograd, "Matrix multiplication via arithmetic progressions," *Journal of Symbolic Computation*, vol. 9, no. 3, pp. 251–280, 1990, computational algebraic complexity editorial.
- [19] H. Kwakernaak and R. Sivan, *Linear Optimal Control Systems*. USA: John Wiley & Sons, Inc., 1972.
- [20] P. Kürschner, "Balanced truncation model order reduction in limited time intervals for large systems," *Advances in Computational Mathematics*, vol. 44, no. 6, pp. 1821–1844, 2018.
- [21] K. Kunisch and S. Volkwein, "Control of the Burgers equation by a reduced-order approach using proper orthogonal decomposition," *Journal of optimization theory and applications*, vol. 102, no. 2, pp. 345–371, 1999.
- [22] J. A. Atwell, J. T. Borggaard, and B. B. King, "Reduced order controllers for Burgers' equation with a nonlinear observer," *International Journal of Applied Mathematics and Computer Science*, vol. 11, pp. 1311–1330, 2001.
- [23] K. Kunisch and L. Xie, "POD-based feedback control of the Burgers equation by solving the evolutionary HJB equation," *Computers & Mathematics with Applications*, vol. 49, no. 7-8, pp. 1113–1126, 2005.
- [24] B. Kramer, "Model reduction for control of a multiphysics system: Coupled Burgers' equation," in *2016 American Control Conference (ACC)*. IEEE, 2016, pp. 6146–6151.
- [25] J. Borggaard and L. Zietsman, "The quadratic-quadratic regulator problem: approximating feedback controls for quadratic-in-state nonlinear systems," in *2020 American Control Conference (ACC)*. IEEE, 2020, pp. 818–823.
- [26] P. Benner, P. Goyal, and S. Gugercin, " \mathcal{H}_2 -quasi-optimal model order reduction for quadratic-bilinear control systems," *SIAM Journal on Matrix Analysis and Applications*, vol. 39, no. 2, pp. 983–1032, 2018.
- [27] A. C. Antoulas, I. V. Gosea, and M. Heinkenschloss, "On the Loewner framework for model reduction of Burgers equation," in *Active Flow and Combustion Control 2018*. Springer, 2019, pp. 255–270.
- [28] B. Kramer and K. Willcox, "Balanced truncation model reduction for lifted nonlinear systems," in *Realization and Model Reduction of Dynamical Systems*. Springer, 2020 (to appear), arXiv:1907.12084.