Model Reduction by Balanced Truncations. Notes by Eduardo Sontag, Rutgers, Nov 2006

This is very summarized. See Antoulas' book for missing proofs (Athanasios C. Antoulas, Approximation of Large-Scale Dynamical Systems, SIAM 2005).

We deal with systems with outputs $\dot{x} = Ax + Bu$, $y = Cx$, and assume that the system is asymptotically stable (A is Hurwitz) and reachable and observable. The goal is to obtain a smaller model that approximates the i/o behavior from the zero initial state.

The idea of the balancing approach is to delete from the model those states that are simultaneously "relatively less observable" and "relatively less reachable (from the origin)". Intuitively, such states should not matter (much) since they will tend not to arise, and when they do arise, they have little influence on the output.

We quantify the degree of reachability and observability by gramians. Recall that

$$W_c == \int_0^\infty e^{tA} BB' e^{tA'}\, dt$$

(this was "W_∞" in earlier notes) has the property that $x' W_c^{-1} x$ is the infimum of the (square of the) possible costs (measured as L^2 norms of inputs) among all possible inputs that drive the origin to the state x in some (arbitrary) finite time. Also, we showed that W_c is a positive definite (symmetric) matrix that satisfies the following matrix equation:

$$AW_c + W_c A' + BB' = 0\,.$$

(Remark: We used that A is Hurwitz in the construction of W_c. A converse also holds: If the pair (A, B) is reachable, and some positive definite matrix W_c satisfies this equation, then A must be Hurwitz; see Exercise 5.7.9 in MCT.)

If $x' W_c^{-1} x$ is large, this means that it is unlikely that this state will appear during normal operation of the system, because very large inputs would have been required in order to attain the state x. In that sense, such a state is "almost unreachable". More generally, components of x along eigendirections corresponding to large eigenvalues of W_c^{-1} will be harder to reach than components in directions corresponding to small eigenvalues, in the following sense.

Arrange the eigenvalues in descending order: $\lambda_1 \geq \ldots \geq \lambda_n$ Suppose that we have found a basis $\{v_1, \ldots, v_n\}$ of $\mathbb{R}^n$ that is orthonormal (i.e., $v_j' v_i = \delta_{ij}$ for all i, j) and so that each v_i is an eigenvector of W_c: $W_c v_i = \lambda_i v_i$.

Multiplying this last equation by $(1/\lambda_i) W_c^{-1}$, we also have that $W_c^{-1} v_i = (1/\lambda_i) v_i$ for each i, i.e., v_i is an eigenvector of W_c^{-1} associated to the eigenvalue $1/\lambda_i$.

It follows that $v_j' W_c^{-1} v_i = v_j' (1/\lambda_i) v_i = \delta_{ij} (1/\lambda_i)$ and also $v_j' W_c v_i = \delta_{ij} \lambda_i$.

Let us assume that, for some $k \in \{1, \ldots, n\}$, there is a gap in eigenvalues: $\lambda_k \gg \lambda_{k+1}$.

Write a state x in this basis: $x = \sum_{i=1}^n a_i v_i$, and let $x_1 = \sum_{i=1}^k a_i v_i$ be the component in the span of $\{v_1, \ldots, v_k\}$ and $x_2 = \sum_{i=k+1}^n a_i v_i$ the component in the span of $\{v_{k+1}, \ldots, v_n\}$. Using orthogonality, we have that

$$x' W_c^{-1} x = \sum_1^n a_i^2 \frac{1}{\lambda_i}\,.$$

Since $1/\lambda_n \geq \ldots \geq 1/\lambda_{k+2} \geq 1/\lambda_{k+1}$, we have that

$$x' W_c^{-1} x \geq \frac{1}{\lambda_{k+1}} \sum_{i=k+1}^{n} a_i^2 \geq \frac{1}{\lambda_{k+1}} \|x_2\|^2 \,.$$

Suppose that the component x_2 of x along $v_{k+1}, \ldots, v_n$ has unit norm. Then, this argument says that the norm of an input that is needed in order to reach x (from the zero initial state) is at least $1/\sqrt{\lambda_{k+1}}$.

On the other hand, if the component x_2 is negligible, $x_2 \approx 0$, then:

$$x' W_c^{-1} x \approx \sum_1^k a_i^2 \frac{1}{\lambda_i} \leq \frac{1}{\lambda_k} \sum_1^k a_i^2 = \frac{1}{\lambda_k} \|x_1\|$$

(since $1/\lambda_1 \leq 1/\lambda_2 \leq \ldots 1/\lambda_k$). If x has unit norm, an input that allows reaching x can be found with norm $\leq 1/\sqrt{\lambda_k}$. As we assumed that $1/\lambda_{k+1} \gg 1/\lambda_k$, we see that a much larger input is needed to reach a state that has a substantial component along $v_{k+1}, \ldots, v_n$ compared to one that does not. Thus, it makes sense to consider only those states in the span of $\{v_1, \ldots, v_k\}$.

A geometric interpretation is as as follows. Consider inputs of unit norm, and let S be the set of those states that can be reached from the origin using such inputs. This is the set defined by $x' W_c^{-1} x \leq 1$. Expressed on the basis $\{v_1, \ldots, v_n\}$, this set is an ellipsoid $\sum_1^n \frac{1}{\lambda_i} x_i^2 \leq 1$, with axes along the spans of the v_i's. The larger axes correspond to small coefficients $\frac{1}{\lambda_i}$, and the smaller axes ("flat" sides) to large coefficients $\frac{1}{\lambda_i}$. (Plot $100x^2 + y^2 = 1$: the ellipsoid is very vertical, flat long the x axis.) Thus, components along the flat sides, i.e. small λ_i's, are relatively small. If the ratio between axes is very large, we can basically think of reachable states as being along the subspace determined by the large axes.

There is an analog for observability. Given any state x, and assuming no inputs are applied, the output that results is the function $y(t) = Ce^{tA}x$. This output has L^2 norm $x' W_o x$, where

$$W_o = \int_0^\infty e^{tA'} C' C e^{tA} \, dt$$

and just as earlier we can prove that W_o satisfies the following matrix equation:

$$A' W_o + W_o A + C' C = 0 \,.$$

What does it mean to say that $x \in \mathbb{R}^n$ is so that $x' W_o x$ is small? It means that the state x is "almost unobservable" since the output that is obtained from it is very small (and possibly undetectable by any measurement devices). An analogous discussion as for reachability suggests that one should reduce a model by dropping components in directions corresponding to eigenvectors associated to small eigenvalues of W_o.

There is a problem, however, with viewing W_c and W_o separately. For example, suppose that we consider the system

$$\begin{aligned} \dot{x}_1 &= -a_1 x_1 + 1000u \\ \dot{x}_2 &= -a_2 x_2 + u \end{aligned}$$

with output $y = 0.001x_1 + x_2$, where $a_1 \neq a_2$ are positive and close to each other. (We pick them different in order to have reachability and observability.) A reachability reduction would tell us to delete the x_2 component, but an observability reduction would tell us to delete the x_1 component. Neither makes sense, since the output for any input $u(\cdot)$ has comparable components in both.

What is needed is that states that are "relatively less observable" be exactly the same ones that are "relatively less reachable". The neat thing about balancing is that this objective can always be achieved, after a change of coordinates.

A change of coordinates is a transformation $z = Tx$, where T is invertible. In the new coordinates, substituting (A, B, C) by the new matrices (TAT^{-1}, TB, CT^{-1}), we get new controllability and observability gramians as follows:

$$\text{new } W_c = \int_0^\infty e^{tTAT^{-1}} TBB'T' e^{t(TAT^{-1})'}\, dt = TW_cT'$$

(so that W_c^{-1} changes to $(T')^{-1}W_c^{-1}T^{-1}$), and

$$\text{new } W_o = \int_0^\infty e^{t(TAT^{-1})'} (T-1)'C'CT-1e^{tTAT^{-1}}\, dt = (T^{-1})'W_oT^{-1}\,.$$

(The product matrix W_oW_c changes as $(T')^{-1}W_oW_cT'$. The square roots of the eigenvalues of this product matrix are called the *Hankel singular values* of the system, and the transformation formula shows that they are not modified under basis changes. The name Hankel singular value arises from the fact that these numbers are the eigenvalues of the "Hankel operator" that maps inputs on $t \le 0$ to outputs for $t \ge 0$.)

Theorem. *There is a set of coordinates in which the system is balanced, that is,* $W_c = W_o$*.*

After this coordinate change in states, we can find a simultaneous basis of orthogonal eigenvectors for both W_c and W_o (since they are the same matrix), and hence the above discussion tells us that states that are "relatively less observable" will be exactly the same ones that are "relatively less reachable," like we wanted.

In fact, one can directly construct a change of coordinates such that, in new coordinates:

$$W_c = W_o = \Sigma = \operatorname{diag}(\sigma_1, \ldots, \sigma_n)$$

with $\sigma_1 \ge \ldots \ge \sigma_n$, and hence the canonical vectors e_i are are orthonormal basis of eigenvectors (with eigenvalues σ_i respectively). (Note that $W_cW_o = \Sigma^2$ in such a basis, so the σ_i are the Hankel singular values.) The proof is very simple.

First pick an invertible matrix R such that

$$W_c = RR'$$

(sort of a "square root" of W_c). This can always be achieved (even, if desired, with an upper triangular R, in which case one calls R a Cholesky factor of W_o). Indeed, we first write $W_c = V\Lambda V'$ where V is orthogonal and Λ is a diagonal matrix with real positive entries (this can be done because W_c is positive definite). Now we let $R = V\Lambda^{1/2}$, where $\Lambda^{1/2}$ is the diagonal matrix whose entries are the square roots of those of L.

Next, consider the matrix $R'W_oR$. Since W_o is positive definite, this is also (symmetric and) positive definite. Thus we can perform an orthogonal decomposition:

$$R'W_oR = U\Sigma^2U'$$

where U is orthogonal and Σ is a diagonal matrix, i.e., $W_o = (R')^{-1}U\Sigma^2U'R^{-1}$.

Let us pick:

$$T \;=\; \Sigma^{1/2}U'R^{-1}\,.$$

Note that $T^{-1} = RU\Sigma^{-1/2}$, $T' = (R^{-1})'U\Sigma^{1/2}$, and $(T^{-1})' = \Sigma^{-1/2}U'R'$. Therefore:

$$TW_cT' = \Sigma^{1/2}U'R^{-1}(RR')(R^{-1})'U\Sigma^{1/2} = \Sigma^{1/2}U'(R^{-1}R)(R'(R')^{-1})U\Sigma^{1/2} = \Sigma^{1/2}U'U\Sigma^{1/2} = \Sigma^{1/2}\Sigma^{1/2} = \Sigma$$

and:

$$(T^{-1})'W_oT^{-1} = \Sigma^{-1/2}U'R'(R')^{-1}U\Sigma^2U'R^{-1}RU\Sigma^{-1/2}\Sigma^{-1/2}\Sigma^2\Sigma^{-1/2} = \Sigma$$

as well.

Finally, we order the eigenvalues using a permutation matrix as a coordinate change. (Both W_c and W_o transform as UWU', if $T = U$ is orthogonal.)

The *balanced truncation* of order k is given as follows. We write the matrices in the new basis, and in partitioned form:

$$A = \begin{pmatrix} A_1 & A_2 \\ A_3 & A_4 \end{pmatrix}, \quad B = \begin{pmatrix} B_1 \\ B_2 \end{pmatrix}, \quad C = (\, C_1 \; C_2 \,)$$

into blocks representing the first k coordinates and the last $n-k$.

The truncated system is obtained by ignoring the last $n-k$ coordinates:

$$(A_1, B_1, C_1)\,.$$

Writing out the equation (recall $W_c = \Sigma$ in the new coordinates):

$$A\Sigma + \Sigma A' + BB' = 0$$

in partitioned form, the $(1,1)$ block tells us that

$$A_1\Sigma_1 + \Sigma_1 A_1' + B_1 B_1' = 0$$

and similarly the observability gramian gives us

$$A_1'\Sigma_1 + \Sigma_1 A + C_1' C_1 = 0\,.$$

Note that Σ_1, the truncation of Σ, is positive definite. Take $V(x) = x'Px$ as a Lyapunov function. Along trajectories of $\dot{x}_1 = A_1 x_1$, we have that $dV/dt = -\|C_1 x\|^2 \leq 0$. Thus, trajectories remain bounded (stability of truncated system).

As observed earlier, if (A, B) is controllable (or if (A, C) is observable), then the reduced system is in fact asymptotically stable (A Hurwitz), as follows from Exercise 5.7.9 in MCT (LaSalle invariance principle).

To fill-in: if $\lambda_k > \lambda_{k+1}$ (that is, if there was really a gap in eigenvalues, the only case in which one would do a truncation to start with!), then it can be proved that controllability and observability hold, so the reduced system is asymptotically stable.

To fill-in: Similarly to what we did for low-rank matrix approximants in L^2 norm, errors for balanced truncation can be given. For example, consider the L^2 induced operator norm (H_∞ norm). Then the difference between the norm of the original system and the approximant is upper bounded by:

$$2\,(\sigma_{k+1} + \ldots + \sigma_n)$$

i.e. twice the sum of the neglected singular values. (This is conservative; the bound is often tighter. For example, when $k = n-1$, the error is exactly σ_n.)

(Discuss in class: H_∞ norms. Also mention Hankel approximation/AAK: same approximation bounds.)