Interface-Preserving $\mathcal{H}_\infty$ Model Reduction of Network Systems

Noam Leiter and Daniel Zelazo *Senior Member, IEEE*

arXiv:submit/8351280 [math.OC] 9 Nov 2026

***Abstract*—We study projection-based reduced-order models that exactly retain specified input and output interfaces. Admissible projections are characterized geometrically: the retained subspace must contain every input and output direction, which fixes the minimum order. The reduction error factors into input-coupling, internal-dynamics, and output-coupling systems, giving a submultiplicative $\mathcal{H}_\infty$ bound; for positively dominated factors the $\mathcal{H}_\infty$ error equals a dc gain computable from linear solves. Strict dissipativity certifies the stability of every orthogonal reduction, and a Hankel singular value bounds the optimum a posteriori. These results motivate sequential graph-coarsening heuristics for networks. On a $65{,}566$-state chip thermal network, interface-preserving coarsening yields sparse, positive lumped models that are more accurate than Kron reduction of the same interface at every order, with far fewer edges; balanced truncation is more accurate but destroys the network structure.**

***Index Terms*—projection-based model reduction, structured model reduction, positive systems, network systems**

I. Introduction

Model reduction provides low-order surrogates for the analysis and control of large-scale systems [1]. Balanced truncation yields a stable reduced model with an a priori $\mathcal{H}_\infty$ error bound [2], [3] and is routine at 10^4–10^5 states with low-rank Lyapunov solvers [4], while Krylov and other projection methods scale further but do not guarantee stability for an arbitrary projection [3]. For network systems, however, the drawback of these methods is structural: the reduced coordinates mix all nodes, so the reduced model loses its graph interpretation and its inputs and outputs no longer correspond to physical nodes.

Structured reductions of networks preserve the graph topology, which encodes stability, controllability, and synchronization properties. Graph partitions have been used to reduce consensus, passive, and dissipative networks [5]–[7], with clusters selected, for example, by the dissimilarity of nodal responses [8]; see [9] for a survey. Clustering of positive directed networks that preserves positivity and steady-state behavior [10] is the closest prior work to the positively dominated setting considered here. Graph-based $\mathcal{H}_\infty$ reduction of multi-agent systems is studied in [11], a priori $\mathcal{H}_2$ and $\mathcal{H}_\infty$ error bounds for clustering of leader–follower systems are derived in [12], and structure-preserving balanced truncation of networked passive systems is developed in [13]. More recently, the edge weights of a clustered consensus network have been optimized for $\mathcal{H}_2$ or $\mathcal{H}_\infty$ error [14], structure-preserving reduction has been extended to interconnected systems [15], and an $\mathcal{H}_2$-optimal Petrov–Galerkin framework preserves stability and dissipativity for every projection basis by weighting the projection with a Lyapunov certificate, without requiring specified input and output nodes to be retained [16].

In this work we require every actuated and measured node to retain its individual identity in the reduced model, a stronger condition than the approximate interface matching permitted by most clustering methods. This is the projection-based counterpart of Kron reduction [17], which keeps a set of boundary nodes and eliminates the interior through a Schur complement. Kron reduction is exact at dc but carries no error guarantee away from dc, and it densifies the reduced graph, since boundary nodes joined through the interior become adjacent. An interface-invariant reduced model instead aggregates the interior into a coarse graph, so interior dynamics remain in averaged form and the error is controlled in $\mathcal{H}_\infty$. Within this family we ask which reduced orders are achievable, how the $\mathcal{H}_\infty$ error can be bounded, and how the remaining projection freedom should be used. The central tool is a product factorization of the error system that separates input coupling, discarded dynamics, and output coupling. For positively dominated systems it yields an exact expression: the $\mathcal{H}_\infty$ error equals a dc gain that can be evaluated without frequency gridding, even at tens of thousands of states.

The main contributions are as follows. (i) We characterize all interface-invariant orthogonal projections geometrically: the retained subspace must span every input and output direction, which establishes the minimum reduced order and the remaining degrees of freedom. (ii) We derive a product representation of the projection-based reduced-order model (PROM) $\mathcal{H}_\infty$ error, with explicit stability hypotheses; for interface-invariant PROMs the coupling factors are strictly proper. (iii) Although the reduction error system is not positive in general, we show that for positively dominated systems [18] its $\mathcal{H}_\infty$ norm reduces to a closed-form steady-state gain, as for internally positive systems [19]; for strictly dissipative systems, a projection-dependent upper bound follows from linear solves, without a frequency sweep. (iv) We express a projection as a sequence of singleton projections, characterize when the resulting triangle-inequality bound is tight, and use it to motivate sequential graph-coarsening heuristics, with the $(r{+}1)$st Hankel singular value as an a posteriori approximation factor. (v) On a $65{,}566$-state processor thermal network with 30 functional-unit interfaces, we compare interface-preserving graph coarsening with Kron reduction, coordinate deletion, and balanced truncation in accuracy, preserved structure, and

N. Leiter is with Prime Minute, Palo Alto, California 94301, (e-mail: noam.leiter@gmail.com).

D. Zelazo is with the Stephen B. Klein Faculty of Aerospace Engineering, Technion–Israel Institute of Technology, Haifa 32000, Israel (e-mail: dzelazo@technion.ac.il). *Corresponding author.*

This research was supported by Grant No. 2285/20 from the Israel Science Foundation.

cost.

Section II formulates the reduction problem and interface invariance. Section III gives the singleton sequence and product factorization. Section IV develops the positively dominated specialization. Section V presents structured sequential heuristics, Section VI presents the chip thermal case study, and Section VII concludes.

Notations: We write $A \succ 0$ for positive-definite, $A \geq 0$ for entrywise nonnegative, and call A *Metzler* if $[A]_{ij} \geq 0$ for all $i \neq j$. For symmetric $A \in \mathbb{R}^{n\times n}$, $\lambda(A) = \{\lambda_k(A)\}_{k=1}^n$ denotes the eigenvalues in ascending order. $[A]_{ij}$ is the (i,j) entry, $\mathrm{diag}(A,B)$ a block-diagonal matrix, and $\|A\|_2$ the spectral norm. The ith standard basis vector of $\mathbb{R}^n$ is e_i^n. For an index set $\mathcal{I} = \{i_1, \ldots, i_m\}$, $e_{\mathcal{I}}^n = [e_{i_1}^n \ \cdots \ e_{i_m}^n]$. $\mathrm{im}\, M$ is the column space of M. $\mathcal{RH}_\infty^{p\times m}$ denotes the set of stable proper real-rational $p \times m$ transfer matrices. A graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ has vertex set $\mathcal{V}$ and edge set $\mathcal{E}$. Nodes u, v are *adjacent* ($u \sim v$) if $uv \in \mathcal{E}$.

II. Model Reduction and Problem Formulation

In this section we review some fundamental notions from the theory of model reduction. We then develop a specialization of the model reduction problem for multi-agent systems. With this framework we will then formally state the problem we aim to solve.

A. The Model Reduction Problem

We consider an LTI system Σ with minimal state-space realization

$$\Sigma : \begin{cases} \dot{x}(t) &= Ax(t) + Bu(t) \\ y(t) &= Cx(t) + Du(t), \end{cases}$$

where $x(t) \in \mathbb{R}^{n_x}$ is the system state, $u(t) \in \mathbb{R}^{n_u}$ is the input, and $y(t) \in \mathbb{R}^{n_y}$ is the output. We denote the state-space realization of a system by the tuple (A, B, C, D). The transfer function matrix (TFM) of the system Σ is given as

$$\hat{\Sigma}(s) = C(sI_{n_x} - A)^{-1}B + D.$$

Hereafter, the notation $\hat{\Sigma}$, without explicit dependence on s, will be used to denote the TFM of a system Σ. A realization $\Sigma := (A, B, C, D)$ is *minimal* if it is controllable and observable. The *order* of a system is its *McMillan degree*, denoted as $\deg(\Sigma)$.

The model reduction problem is to find an order-r system Σ_r, with $r < n_x$, whose transfer matrix approximates $\hat{\Sigma}$. The reduced and full models retain the same external input and output dimensions: $B_r \in \mathbb{R}^{r\times n_u}$ and $C_r \in \mathbb{R}^{n_y\times r}$. We define the error system

$$\Sigma_e = \Sigma - \Sigma_r, \tag{1}$$

which has realization

$$\begin{aligned} \Sigma_e \ &:= (A_e, B_e, C_e, D_e) \\ &= \left(\begin{bmatrix} A & 0 \\ 0 & A_r \end{bmatrix}, \begin{bmatrix} B \\ B_r \end{bmatrix}, \begin{bmatrix} C & -C_r \end{bmatrix}, D - D_r \right). \end{aligned} \tag{2}$$

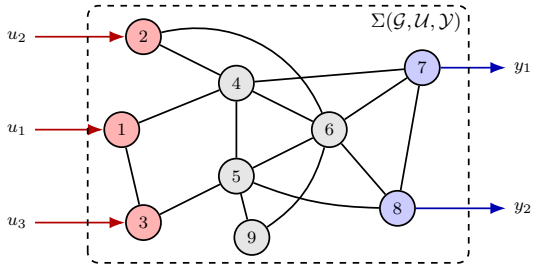


Fig. 1. An MAS with control nodes (red) and output nodes (blue).

Thus, one can consider, for a fixed r, finding a system Σ_r such that $\|\Sigma_e\|$ is minimized. For a detailed review see [3]. In this work we focus on $\mathcal{H}_\infty$ reduction errors. For a stable and proper system,

$$\|\hat{\Sigma}\|_\infty = \sup_{\omega\in\mathbb{R}} \bar{\sigma}(\hat{\Sigma}(j\omega)),$$

where $\bar{\sigma}(M)$ is the largest singular value of the matrix M.

A widely used family of reduction methods are the so-called *projection-based reductions*. A Petrov–Galerkin reduction uses trial and test bases $V, W \in \mathbb{R}^{n_x\times r}$ satisfying $W^TV = I_r$ and gives

$$\Sigma_r(W, V) := (W^TAV, W^TB, CV, D). \tag{3}$$

When $W = V = P$ and $P^TP = I_r$, the reduction is orthogonal. We denote the resulting realization by $\Sigma_r(P)$; it is obtained from (3) by setting $W = V = P$. We call this model an *orthogonal projected reduced-order model* (PROM).

B. Multi-Agent Systems and Interface-Invariant Model Reduction

We consider linear multi-agent systems (MAS) comprised of n agents interacting over a network $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, with state $x_i \in \mathbb{R}^{q_x}$ per agent ($nq_x = n_x$). The block structure of $A(\mathcal{G})$ follows the graph: the ij-block is nonzero iff $ij \in \mathcal{E}$. Input nodes $\mathcal{U} \subseteq \mathcal{V}$ ($|\mathcal{U}| = m$) receive external inputs $u_i \in \mathbb{R}^{q_u}$; output nodes $\mathcal{Y} \subseteq \mathcal{V}$ ($|\mathcal{Y}| = p$) are measured, $y_i \in \mathbb{R}^{q_y}$. The MAS realization is

$$\Sigma(\mathcal{G},\mathcal{U},\mathcal{Y}) := (A(\mathcal{G}), B(\mathcal{U}), C(\mathcal{Y}), D(\mathcal{U},\mathcal{Y})), \tag{4}$$

with $B(\mathcal{U}) \in \mathbb{R}^{n_x\times n_u}$, $n_u = mq_u$, and $C(\mathcal{Y}) \in \mathbb{R}^{n_y\times n_x}$, $n_y = pq_y$ (Figure 1). Agents in $\mathcal{U}\cup\mathcal{Y}$ are the *interface agents*; self-loops are permitted.

In this work we will examine a special class of PROMs that we term *interface-invariant PROMs*. These are reduced models that maintain the input-output structure of the full order model under the projection operation. Such PROMs are required, for example, for the reduction of the MAS (4) where the interface agent structure is maintained in the reduced model. We now define formally the notion of an interface-invariant PROM (IIPROM).

Definition 1 (IIPROM)**.** *Given a system with realization $\Sigma := (A, B, C, D)$, an* interface invariant projection-based reduced order model *(IIPROM) of Σ is any PROM $\Sigma_r(P)$ such that $C = CPP^T$ and $B = PP^TB$.*

Note that an IIPROM does not require $PP^T = I_{n_x}$. For example, if $C = B^T = [1\ 0\ 0]$, the orthonormal matrix

$$P = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1/\sqrt{2} & 1/\sqrt{2} \end{bmatrix}^T$$

satisfies $C = CPP^T$ and $B = PP^TB$.

Remark 1. *Although motivated by multi-agent systems, the IIPROM framework is not intrinsically graph based. It applies to any orthogonal PROM for which prescribed input and output state-space directions must be retained, regardless of whether the system has a network interpretation.*

Define the *protected input–output subspace*

$$\mathcal{S}_{\mathrm{io}} := \operatorname{im} B + \operatorname{im} C^T = \operatorname{im} \begin{bmatrix} B & C^T \end{bmatrix}, \quad k := \dim \mathcal{S}_{\mathrm{io}}.$$

The following result separates the interface requirement from the remaining model-reduction freedom.

Theorem 1 (Geometric characterization of IIPROMs)**.** *Let $P \in \mathbb{R}^{n_x \times r}$ satisfy $P^TP = I_r$, and let $\Pi = PP^T$. Then $\Sigma_r(P)$ is interface invariant if and only if*

$$\mathcal{S}_{\mathrm{io}} \subseteq \operatorname{im} P. \tag{5}$$

Consequently, an order-r interface-invariant projection exists if and only if $k = \dim \mathcal{S}_{\mathrm{io}} \le r \le n_x$, and it is a strict reduction precisely when $k \le r < n_x$.

More precisely, let $U_{\mathrm{io}} \in \mathbb{R}^{n_x \times k}$ be an orthonormal basis for $\mathcal{S}_{\mathrm{io}}$, and complete it to the orthogonal matrix $[\,U_{\mathrm{io}}\ U_\perp\,]$. Every admissible retained subspace has an orthonormal basis P of the form

$$P = \begin{bmatrix} U_{\mathrm{io}} & U_\perp Z \end{bmatrix} R,\ Z^TZ = I_{r-k},\ R^TR = RR^T = I_r, \tag{6}$$

where $Z \in \mathbb{R}^{(n_x-k)\times(r-k)}$ and $R \in \mathbb{R}^{r\times r}$ is an orthogonal basis-change within the retained subspace; since $\Pi = PP^T$ is invariant under right-multiplication by R, the factor R does not affect the reduced model. Conversely, every matrix in (6) *defines an order-r IIPROM.*

Proof. Because Π is the orthogonal projector onto $\operatorname{im} P$, $\Pi B = B$ holds if and only if $\operatorname{im} B \subseteq \operatorname{im} P$ (similarly condition for $\Pi C^T = C^T$). The two interface equalities are therefore equivalent to (5), which also proves the order condition.

For the parameterization, suppose (5) holds. The retained subspace has the orthogonal decomposition

$$\operatorname{im} P = \mathcal{S}_{\mathrm{io}} \oplus \mathcal{W}, \qquad \mathcal{W} \subseteq \mathcal{S}_{\mathrm{io}}^{\perp}, \qquad \dim \mathcal{W} = r - k.$$

Writing an orthonormal basis of $\mathcal{W}$ in the coordinates of $U_\perp$ gives $U_\perp Z$ with $Z^TZ = I_{r-k}$. Thus $P_0 = [\,U_{\mathrm{io}}\ U_\perp Z\,]$ and P are orthonormal bases for the same subspace, so $P = P_0R$ for an orthogonal matrix R. The converse follows because every such P contains $\mathcal{S}_{\mathrm{io}}$ in its range. □

Right multiplication by R only changes reduced coordinates. The matrices $\Sigma_r(P_0R)$ and $\Sigma_r(P_0)$ are similar and have the same transfer matrix. Hence one may set $R = I_r$ without loss of input–output generality. The constrained search over P is thereby reduced to the choice of an $(r-k)$-dimensional subspace of $\mathcal{S}_{\mathrm{io}}^{\perp}$, represented by Z in (6). For a multi-agent system with coordinate-selecting B and C, $\mathcal{S}_{\mathrm{io}}$ is simply the span of the actuated and measured agent coordinates.

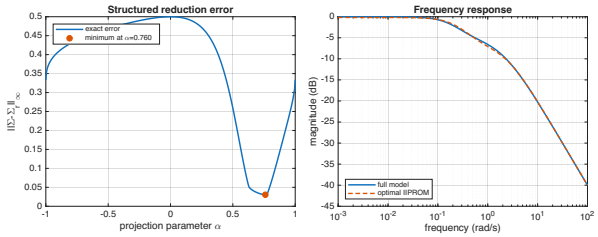


Fig. 2. Example 1: reduction error $\|\Sigma_e\|_\infty$ vs. α with its minimizer (left), and the full and optimal reduced frequency responses (right).

With the above notions in place, we can now formally state the optimal IIPROM problem.

Problem 1 (optimal IIPROM)**.** *Consider a stable proper system of order n_x with realization $\Sigma := (A, B, C, D)$. Find $P \in \mathbb{R}^{n_x\times r}$ with $P^TP = I_r$ such that the PROM $\Sigma_r(P)$ is interface invariant and stable and minimizes the $\mathcal{H}_\infty$-norm of the reduction error system* (1)*,*

$$\begin{aligned} \min_{P\in\mathbb{R}^{n_x\times r}} & \|\Sigma_e\|_\infty \\ \text{s.t.} \quad & P^TP = I_r, \quad C = CPP^T, \\ & B = PP^TB, \quad P^TAP \text{ Hurwitz.} \end{aligned}$$

Theorem 1 removes the interface-invariance constraints from Problem 1 by replacing P with the parameterization (6), reducing the search to the choice of $Z \in \mathbb{R}^{(n_x-k)\times(r-k)}$ with orthonormal columns. The remaining $\mathcal{H}_\infty$ objective is nevertheless nonsmooth and nonconvex, and reduced-model stability remains a nonlinear constraint. Consequently, the problem generally calls for optimization on the reduced Grassmann manifold associated with $\mathcal{S}_{\mathrm{io}}^{\perp}$ or for a structured restriction of the feasible projections; a Riemannian method of this kind for structure-preserving $\mathcal{H}_2$ reduction is developed in [16].

While a closed-form solution is not available in general, there are instances where an analytic solution exists. We provide a simple example below.

Example 1. *Consider the SISO system $\Sigma = (A, B, C, D)$ with $A = \left[\begin{smallmatrix} -2 & 1 & 0 \\ 1 & -2 & 1 \\ 0 & 1 & -1 \end{smallmatrix}\right]$, $B = C^T = e_1$, and $D = 0$, whose transfer function is $\hat{\Sigma} = (s^2 + 3s + 1)/(s^3 + 5s^2 + 6s + 1)$, with $\|\Sigma\|_\infty = 1$. We seek an IIPROM of order $r = 2$. Every $P \in \mathbb{R}^{3\times 2}$ with $P^TP = I_2$, $CPP^T = C$, and $PP^TB = B$ has the form $P(\alpha) = [e_1 \quad \alpha e_2 + \beta(\alpha)e_3]$, with $\alpha \in [-1, 1]$ and $\beta(\alpha) = \sqrt{1-\alpha^2}$. The resulting IIPROM is*

$$\hat{\Sigma}_r = \frac{s + \gamma^2}{s^2 + (\gamma^2 + 2)s + 2\gamma^2 - \alpha^2},$$

with $\gamma^2 = \alpha^2 + (\alpha - \beta(\alpha))^2$. It is stable for all α and minimal for $\alpha \neq 0$; at $\alpha = 0$ the retained direction e_3 is decoupled from e_1 and $\hat{\Sigma}_r = 1/(s + 2)$. The error $\|\Sigma_e\|_\infty$ ranges from 0.5 at $\alpha = 0$ to its minimum ≈ 0.03 at $\alpha^* = 0.76$, which solves Problem 1, and it is nonconvex and nonsmooth in α (Fig. 2).

Example 1 exploits the fact that the free subspace in Theorem 1 is one dimensional. In general the geometric parameterization removes the two interface equalities but does not produce a closed-form or convex solution of Problem 1: the choice of the complementary subspace, the $\mathcal{H}_\infty$ objective, and the reduced stability condition remain coupled. In the sequel, we investigate the IIPROM error structure and derive tractable suboptimal constructions.

III. Structure and Analysis of PROMs and IIPROMs

Understanding the structure of IIPROMs is essential for tackling Problem 1 for multi-agent systems. In this direction, we provide complementary perspectives on the error system associated with IIPROMs. First, we introduce a sequential decomposition of the projection matrix used in PROMs into a product of elementary "singleton" projections, each reducing the system order by one. This construction leads to a telescoping sum representation of the error system and yields a tractable upper bound on the overall reduction error. Building on this result, we formulate a suboptimal IIPROM design problem aimed at minimizing this bound. These results provide a foundation for algorithmic strategies that sequentially eliminate state variables while preserving input-output invariance and stability.

A. The PROM Sequence Bound

Stability of the full-order model alone does not guarantee stability of an orthogonal PROM. The following proposition gives two sufficient conditions under which an orthogonal compression is stable. It will be used both for the singleton sequence below and later for graph-contraction candidates.

Proposition 1 (A priori stability of orthogonal compressions)**.** *Let $P^TP = I_r$, $\Pi = PP^T$, and $A_P = P^TAP$.*

1) *If, for some $\mu > 0$,*
$$\frac{A + A^T}{2} \preceq -\mu I_n, \tag{7}$$
then A_P is Hurwitz for every P, with $\max \mathrm{Re}\,\lambda(A_P) \leq -\mu$.
2) *More generally, suppose $X = X^T \succ 0$ satisfies $A^TX + XA \prec 0$. If $X\Pi = \Pi X$, then A_P is Hurwitz with Lyapunov matrix $X_P = P^TXP \succ 0$.*

Proof. For the first claim,
$$\frac{A_P + A_P^T}{2} = P^T \frac{A + A^T}{2} P \preceq -\mu I_r,$$
which bounds the real part of every eigenvalue of A_P. For the second claim, commutation gives $XP = PX_P$ and $P^TX\Pi = P^TX$. Hence
$$A_P^TX_P + X_PA_P = P^T(A^TX + XA)P \prec 0.$$
□

In particular, under part (i), every intermediate compression in any singleton projection sequence is Hurwitz with the same decay margin μ. More generally, part (ii) applies whenever the chosen retained subspace is compatible with a Lyapunov matrix through the commutation condition. Some additional assumption beyond mere Hurwitz stability of A is essential, since a Hurwitz nonsymmetric matrix can have an unstable orthogonal compression. The commutation condition can be avoided by replacing the orthogonal projection with the X-orthogonal Petrov–Galerkin projection $P(P^TXP)^{-1}P^TX$, which preserves stability, and more generally any dissipativity certified by X, for every P [16]. We retain orthogonal projections because the product form of Section III-B and the quotient-graph interpretation of the partition projections in Section V rely on them.

In the following section we present a lemma showing that any projection from $\mathbb{R}^{n_x}$ to $\mathbb{R}^r$ can be obtained from a sequence of $n_x - r$ projections, each reducing the dimension by one. We denote such projections as *singleton projections*.

Lemma 1 (Singleton factorization)**.** *Let $P \in \mathbb{R}^{n_x \times r}$ with $P^TP = I_r$ and $r < n_x$. Then there exists a sequence $\{P_{(k)}\}_{k=1}^{n_x-r}$ with $P_{(k)} \in \mathbb{R}^{n_x-k+1 \times n_x-k}$ and $P_{(k)}^TP_{(k)} = I_{n_x-k}$, such that $P = \prod_{k=1}^{n_x-r} P_{(k)}$ (with factors ordered from left to right in increasing k).*

Proof. Let $P \in \mathbb{R}^{n_x \times r}$ with $P^TP = I_r$. Then there exists $Q \in \mathbb{R}^{n_x \times n_x-r}$ with $Q^TQ = I_{n_x-r}$ such that $P^TQ = 0$. Construct
$$P_{(1)} = \begin{bmatrix} P & q_1 & \cdots & q_{n_x-r-1} \end{bmatrix},$$
where q_i is the ith column of Q, and
$$P_{(k)} = \begin{bmatrix} I_{n_x-k} & 0_{n_x-k \times 1} \end{bmatrix}^T$$
for $k \in [2,\, n_x - r]$. Since $P^TP = I_r$ and $P^TQ = 0$ we have $P_{(1)}^TP_{(1)} = I_{n_x-1}$, and for $k \in [2,\, n_x - r]$ it follows that $P_{(k)}^TP_{(k)} = I_{n_x-k}$. From this construction we get
$$\begin{aligned}\prod_{k=1}^{n_x-r} P_{(k)} &= P_{(1)} \left(\prod_{i=1}^{n_x-r-1} P_{(1+i)} \right) \\ &= P_{(1)} \left(\prod_{i=1}^{n_x-r-1} \begin{bmatrix} I_{n_x-i} & 0_{(n_x-i)\times 1} \end{bmatrix}^T \right) = P.\end{aligned}$$
□

Note that one can generate other sequences $P = \prod_{k=1}^{n_x-r} \tilde{P}_{(k)}$ using the transformations $\tilde{P}_{(1)} = P_{(1)}U_{(1)}$, $\tilde{P}_{(k)} = U_{(k-1)}^TP_{(k)}U_{(k)}$ for $k \in [2,\, n_x - r - 1]$, and $\tilde{P}_{(n_x-r)} = U_{(n_x-r-1)}^TP_{(n_x-r)}$, where $U_{(k)} \in \mathbb{R}^{(n_x-k)\times(n_x-k)}$ is an orthogonal matrix.

By expressing a PROM as a sequence of singleton projections, we obtain the following PROM sequence bound.

Proposition 2 (Singleton PROM sequence bound)**.** *Consider an LTI system $\Sigma_{(0)} := \big(A_{(0)}, B_{(0)}, C_{(0)}, D\big)$ and a PROM $\Sigma_r(P)$. Let $\big\{P_{(k)}\big\}_{k=1}^{n_x-r}$ be a sequence with $P_{(k)} \in \mathbb{R}^{(n_x-k+1)\times(n_x-k)}$, $P_{(k)}^TP_{(k)} = I_{n_x-k}$ such that $P = \prod_{k=1}^{n_x-r} P_{(k)}$. Define recursively*
$$\begin{aligned} &A_{(k)} = P_{(k)}^TA_{(k-1)}P_{(k)}, \qquad B_{(k)} = P_{(k)}^TB_{(k-1)}, \\ &C_{(k)} = C_{(k-1)}P_{(k)}. \end{aligned}$$

Assume every $A_{(k)}$, including $A_{(0)}$, is Hurwitz. Then the $\mathcal{H}_\infty$ norm of the error reduction system, Σ_e in (1)*, is bounded by*

$$\|\Sigma_e\|_\infty \le \sum_{k=1}^{n_x-r} \|\Sigma_{(k)} - \Sigma_{(k-1)}\|_\infty, \tag{8}$$

where $\Sigma_{(k)} := (A_{(k)}, B_{(k)}, C_{(k)}, D)$.

Proof. We express $\Sigma_e = \Sigma_r - \Sigma$ as the telescoping sum

$$\sum_{k=1}^{n_x-r} \left(\Sigma_{(k)} - \Sigma_{(k-1)}\right),$$

with $\Sigma_{(0)} = \Sigma$ and $\Sigma_{(n_x-r)} = \Sigma_r(P)$. From the triangle inequality we obtain

$$\|\Sigma_e\|_\infty \le \sum_{k=1}^{n_x-r} \|\Sigma_{(k)} - \Sigma_{(k-1)}\|_\infty,$$

completing the proof. □

The bound (8) can be conservative when the step errors peak at different frequencies or input directions. The following result identifies the precise conditions under which equality holds.

Proposition 3 (Tightness of the singleton sequence bound)**.** *For $k = 1, \dots, n_x - r$, define the stable, strictly proper step error*

$$E_k(s) := \hat{\Sigma}_{(k-1)}(s) - \hat{\Sigma}_{(k)}(s), \qquad a_k := \|E_k\|_\infty,$$

and omit any identically zero E_k. Equality holds in (8) *if and only if there exist a frequency $\omega_\star \in \mathbb{R}$ and unit vectors $v_\star \in \mathbb{C}^{n_u}$ and $u_\star \in \mathbb{C}^{n_y}$ such that*

$$E_k(j\omega_\star)v_\star = a_k u_\star \quad \textit{for every nonzero step } k. \tag{9}$$

Thus all steps must attain their $\mathcal{H}_\infty$ norms at one frequency, on the same worst-case input direction, with phase-aligned output directions.

Proof. The telescoping identity gives $\hat{\Sigma} - \hat{\Sigma}_r = \sum_k E_k$. If (9) holds, then

$$\left\| \sum_k E_k(j\omega_\star)v_\star \right\|_2 = \sum_k a_k,$$

so the sequence bound is attained. Conversely, the nonzero strictly proper rational matrix $\sum_k E_k$ attains its $\mathcal{H}_\infty$ norm at a finite frequency. If the sequence bound is attained, choose a unit maximizing right singular vector $v_\star$ at such a frequency $\omega_\star$. Then every inequality in

$$\begin{aligned}\left\| \sum_k E_k(j\omega_\star)v_\star \right\|_2 &\le \sum_k \|E_k(j\omega_\star)v_\star\|_2 \\ &\le \sum_k \bar{\sigma}(E_k(j\omega_\star)) \le \sum_k a_k\end{aligned}$$

must be an equality. Equality in the Euclidean triangle inequality makes the vectors $E_k(j\omega_\star)v_\star$ nonnegative multiples of a common unit vector $u_\star$, while equality in the remaining inequalities makes each multiple equal to a_k. This is (9). □

When the step errors are positively dominated, the peak frequency of each E_k is pinned to dc and the alignment condition of Proposition 3 reduces to a purely algebraic criterion on the dc gains; this is explored in detail in Section IV.

Proposition 3 separates two sources of conservatism in the sequence bound. Different peak frequencies prevent frequency alignment, while different singular directions cause cancellation or dispersion even when all steps peak at the same frequency. The positively dominated specialization developed in Section IV removes the first source.

With Proposition 2, we can consider a suboptimal variation of Problem 1 that attempts to minimize the upper bound for the reduction error.

Problem 2 (suboptimal IIPROM sequence bound)**.** *Consider a stable proper system Σ of order n_x, with realization (A, B, C, D). Find a sequence $\{P_{(k)}\}_{k=1}^{n_x-r}$ such that the associated PROM is interface invariant, all intermediate PROMs are stable, and the sequence bound* (8) *is minimized:*

$$\begin{aligned}\min_{\{P_{(k)}\}} \ & \sum_{k=1}^{n_x-r} \|\Sigma_{(k)} - \Sigma_{(k-1)}\|_\infty \\ s.t.\ \ & P = \prod_{k=1}^{n_x-r} P_{(k)},\ P_{(k)} \in \mathbb{R}^{(n_x-k+1)\times(n_x-k)}, \\ & P_{(k)}^T P_{(k)} = I_{n_x-k}, \quad A_{(k)} \textit{ Hurwitz for all } k, \\ & C = CPP^T,\ B = PP^T B,\end{aligned} \tag{10}$$

where $\Sigma_{(k)}$ is defined recursively as in Proposition 2.

Problem 2 is not a convex relaxation of Problem 1. Its practical benefit is algorithmic: it permits a greedy search over admissible singleton reductions, with stability and interface invariance checked after every step. The accumulated objective is a certified upper bound by the triangle inequality, but its conservatism can grow with the number of steps.

B. The Product Form of Orthogonal PROMs

It will also be useful to represent a PROM as a subsystem of Σ after an appropriately defined transformation. We apply the unitary transformation $U = \begin{bmatrix} P & Q \end{bmatrix}$ on Σ, where Q is the orthogonal complement of P, satisfying $Q^T P = 0$ and $Q^T Q = I_{n_x-r}$. The resulting realization is

$$\tilde{\Sigma} := \left(U^T AU, U^T B, CU, D\right), \tag{11}$$

which can be expanded to

$$\tilde{\Sigma} := \left(\left[\begin{array}{c|c} A_{PP} & A_{PQ} \\ \hline A_{QP} & A_{QQ} \end{array}\right], \left[\begin{array}{c} P^T B \\ \hline Q^T B \end{array}\right], \left[\begin{array}{c|c} CP & CQ \end{array}\right], D \right), \tag{12}$$

where $A_{PP} \triangleq P^T AP$, $A_{PQ} \triangleq P^T AQ$, $A_{QP} \triangleq Q^T AP$ and $A_{QQ} \triangleq Q^T AQ$.

Remark 2 (Interface invariance in PQ coordinates)**.** *If Σ_r is an IIPROM, then $Q^T B = 0$ and $CQ = 0$ in* (12)*, since $Q^T B = Q^T(PP^T B) = 0$ and $CQ = (CPP^T)Q = 0$. Consequently, the PQ realization of any IIPROM is block-lower-triangular: the retained-subspace dynamics are decoupled from the discarded inputs and outputs.*

The sequence bound of the previous subsection controls the cumulative error but treats each step error $\|\Sigma_{(k)} - \Sigma_{(k-1)}\|_\infty$ as an opaque quantity. To sharpen the analysis, we need to

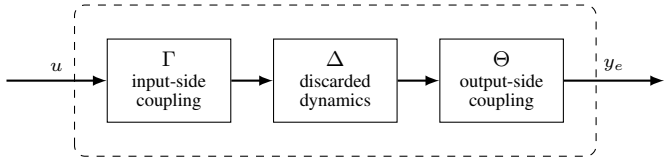


Fig. 3. Product (cascade) representation of the PROM error system $\Sigma_e = \Theta(\Sigma, P)\,\Delta(\Sigma, P)\,\Gamma(\Sigma, P)$, separating input-side coupling, discarded dynamics, and output-side coupling.

understand the internal structure of the PROM error system. Specifically, how the input coupling, the discarded dynamics, and the output coupling each contribute. We show that the error system of any orthogonal PROM factors as a cascade of three LTI systems that isolate these three contributions, reducing the full-order augmented realization of dimension $\deg(\Sigma) + \deg(\Sigma_r)$ to a product of systems of dimensions $\deg(\Sigma_r)$, $\deg(\Sigma)$, and $\deg(\Sigma_r)$.

Theorem 2. *Let $\Sigma := (A, B, C, D)$, and consider a PROM $\Sigma_r(P)$. At every s for which the inverses below exist, the error system $\Sigma_e = \Sigma - \Sigma_r(P)$ has the TFM*

$$\hat{\Sigma}_e = C\Phi^{-1}Q(Q^T\Phi^{-1}Q)^{-1}Q^T\Phi^{-1}B, \tag{13}$$

where $\Phi \triangleq sI_{n_x} - A$, and Q satisfies $Q^TP = 0$ and $Q^TQ = I_{n_x-r}$. Furthermore, Σ_e can be expressed as the product of three systems,

$$\Sigma_e = \Theta(\Sigma, P)\Delta(\Sigma, P)\Gamma(\Sigma, P), \tag{14}$$

with realizations

$$\Theta(\Sigma, P) := (A_{PP}, A_{PQ}, CP, CQ), \tag{15a}$$
$$\Gamma(\Sigma, P) := (A_{PP}, P^TB, A_{QP}, Q^TB), \tag{15b}$$
$$\Delta(\Sigma, P) := (A, Q, Q^T, 0_{(n_x-r)\times(n_x-r)}). \tag{15c}$$

The proof of Theorem 2 is given in the Appendix. Because both sides are rational matrices, equality away from their finite pole sets establishes the transfer-matrix identity.

We observe that each of the three systems of the PROM product form (14) implicitly captures the contribution to the reduction error of the parts of $\tilde{\Sigma}$ in (12) left out in $\Sigma_r(\Sigma, P)$: the map $\Theta(\Sigma, P)$ captures A_{PQ} and CQ, the map $\Gamma(\Sigma, P)$ captures A_{QP} and Q^TB and the map $\Delta(\Sigma, P)$ captures the full-order dynamics viewed through the discarded subspace.

Corollary 1 (Frequency-uniform product bound)**.** *If A and A_{PP} are Hurwitz, then all three factors in* (15) *are stable and*

$$\|\Sigma_e\|_\infty \le \underbrace{\|\Theta(\Sigma, P)\|_\infty \|\Delta(\Sigma, P)\|_\infty \|\Gamma(\Sigma, P)\|_\infty}_{:=b_{\mathrm{prod}}(\Sigma,P)}. \tag{16}$$

Proof. The state matrices of the three factors are A_{PP}, A, and A_{PP}. The result follows from submultiplicativity of the induced $\mathcal{L}_2$ norm. □

The certificate in (16) is uniform over frequency. In fact, for every $\omega \in \mathbb{R}$,

$$\begin{aligned}\bar{\sigma}\left(\hat{\Sigma}_e(j\omega)\right) &\le \bar{\sigma}\left(\hat{\Theta}(j\omega)\right)\bar{\sigma}\left(\hat{\Delta}(j\omega)\right)\bar{\sigma}\left(\hat{\Gamma}(j\omega)\right)\\ &\le b_{\mathrm{prod}}(\Sigma, P).\end{aligned}$$

Equivalently, under zero initial conditions, every $u \in \mathcal{L}_2$ satisfies the input–output uncertainty bound

$$\|y - y_r\|_{\mathcal{L}_2} \le b_{\mathrm{prod}}(\Sigma, P)\,\|u\|_{\mathcal{L}_2}.$$

Thus a single scalar controls the error system for all frequencies and all finite-energy inputs. The condition on A_{PP} is essential, since a Hurwitz nonsymmetric A can have an unstable orthogonal compression. Moreover, this is a certificate for a chosen projection, not an a priori bound that is uniform over all admissible IIPROMs or depends only on the reduced order. Without a common stability margin, admissible reduced poles can approach the imaginary axis and the factor norms can become arbitrarily large. A uniform bound over a family of projections therefore requires common stability and factor-norm bounds. Finally, evaluating b_{prod} requires two reduced-order norm computations and one full-order norm computation, so it is not presented as a universal complexity reduction. It instead decomposes the error by mechanism and supplies a useful score when the factor norms or their dc gains are inexpensive to evaluate.

A first observation related to the product form of Theorem 2 is that if the PROM is an IIPROM, then the realizations of $\Theta(\Sigma, P)$ and $\Gamma(\Sigma, P)$ are strictly-proper.

Corollary 2 (IIPROM error system)**.** *If $\Sigma_r(P)$ is an IIPROM, then $\Theta(\Sigma, P)$ in* (15a) *and $\Gamma(\Sigma, P)$ in* (15b) *are strictly proper with realizations*

$$\Theta(\Sigma, P) := \left(A_{PP}, A_{PQ}, CP, 0_{n_y\times(n_x-r)}\right), \tag{17}$$

and

$$\Gamma(\Sigma, P) := \left(A_{PP}, P^TB, A_{QP}, 0_{(n_x-r)\times n_u}\right). \tag{18}$$

Proof. Since $\Sigma_r(P)$ is an IIPROM, from Definition 1 we have $C = CPP^T$ and $B = PP^TB$. From $Q^TP = 0$ and $Q^TQ = I_{n_x-r}$ we get $PP^T = I_{n_x} - QQ^T$ such that $C = C\left(I_{n_x} - QQ^T\right)$ and $B = \left(I_{n_x} - QQ^T\right)B$, therefore, $CQ = 0$ and $Q^TB = 0$, obtaining the stated zero direct terms. □

Interface invariance removes an instantaneous error path through the discarded coordinates, so the error transfer matrix is strictly proper—there is no instantaneous feedthrough from input to error output.

IV. Exact Error for Positively Dominated Factors

The product bound (14) is tight for an important special class of systems—the *positively dominated* systems introduced by Rantzer [18]. We next recall the relevant definitions before deriving the main exactness result.

Positive LTI systems have the unique property of maintaining output within an invariant subspace $\mathbb{R}^p_+$ for non-negative input signals. They are formally defined below.

Definition 2 ([20])**.** *An LTI system $\Sigma := (A, B, C, D)$ is* internal positive *if for every non-negative initial condition $x(0)$ and every non-negative input $u(t)$, $t \ge 0$, the associated state and output evolutions $x(t)$ and $y(t)$, $t \ge 0$, remain non-negative at every time instant t.*

Positive systems admit a special realization.

Proposition 4 ([20]). *A system Σ with realization (A, B, C, D) is* internally positive *if and only if A is Metzler and $B, C, D \geq 0$.*

Rantzer introduced *positively dominated* systems as a generalization of positive systems [18].

Definition 3 (Positively dominated system). *A system $\hat{\Sigma}(s) \in \mathcal{RH}_\infty^{n\times m}$ is called* positively dominated *if every matrix entry satisfies $\left|\hat{\Sigma}_{kl}(j\omega)\right| \leq \hat{\Sigma}_{kl}(0)$ for $\omega \in \mathbb{R}$. The set of all such matrices is denoted $P\mathcal{H}_\infty^{n\times m}$.*

Positively dominated systems have their $\mathcal{H}_\infty$ norm at dc and are closed under compatible products. For internally positive systems the first property is the familiar fact that the $\mathcal{H}_\infty$ norm equals the dc gain [19].

Proposition 5 ([18]). *Let $\hat{\Sigma} \in P\mathcal{H}_\infty^{n\times m}$, then $\|\Sigma\|_{\mathcal{H}_\infty} = \|\hat{\Sigma}(0)\|_2$. Furthermore, for $\hat{\Sigma}_1 \in P\mathcal{H}_\infty^{n\times m}$ and $\hat{\Sigma}_2 \in P\mathcal{H}_\infty^{l\times n}$, then $\hat{\Sigma}_2\hat{\Sigma}_1 \in P\mathcal{H}_\infty^{l\times m}$.*

The tightness condition of Proposition 3 takes a particularly simple form for positively dominated step errors: the peak frequency is pinned to dc, reducing the alignment condition to an algebraic criterion on the dc gain matrices.

Corollary 3 (Positive collocated steps). *Suppose every step error E_k is positively dominated and $E_k(0)$ is symmetric positive semidefinite. Then the sequence bound* (8) *is tight if and only if the dominant eigenspaces of all nonzero $E_k(0)$ share a common unit vector. In particular, if*

$$E_k(0) = \frac{z_k z_k^T}{d_k}, \qquad d_k > 0, \tag{19}$$

then equality holds if and only if all nonzero z_k are collinear. The condition is automatic for scalar input and output.

Proof. Positive dominance pins the common peak frequency to $\omega_\star = 0$. For a positive semidefinite dc matrix, the spectral norm is its largest eigenvalue, so (9) is equivalent to a common dominant eigenvector. The dominant eigenspace of the rank-one matrix in (19) is $\operatorname{span}(z_k)$. □

Remark 3 (Order invariance for rank-one positive steps). *Under the rank-one hypotheses of Corollary 3, positive dominance gives $\|E_k\|_\infty = \|E_k(0)\|_2 = \operatorname{trace} E_k(0)$ for each step. Summing over k, the right-hand side of* (8)*, $U_r := \sum_k \|E_k\|_\infty$, equals $\operatorname{trace} E(0)$, where $E(0) = \sum_k E_k(0)$ is the final dc error—independent of deletion order. Consequently, reordering cannot improve the bound, and it is tight if and only if $\operatorname{rank} E(0) = 1$.*

Theorem 3 (Exact error at dc). *Let A and A_{PP} be Hurwitz. If the three factors $\Theta(\Sigma, P)$, $\Delta(\Sigma, P)$, and $\Gamma(\Sigma, P)$ are positively dominated, then Σ_e is positively dominated and*

$$\|\Sigma_e\|_\infty = \|\hat{\Sigma}_e(0)\|_2 = \left\|CA^{-1}Q(Q^TA^{-1}Q)^{-1}Q^TA^{-1}B\right\|_2. \tag{20}$$

Proof. Closure under multiplication and Theorem 2 imply that Σ_e is positively dominated. Proposition 5 therefore gives $\|\Sigma_e\|_\infty = \|\hat{\Sigma}_e(0)\|_2$. Substitution of $s = 0$ in (13) gives (20); the four minus signs cancel. □

The hypotheses concern the factors, not merely the original realization. Positivity of Σ alone does not guarantee positivity of the factors for an arbitrary dense orthogonal basis because P and Q may have mixed signs. A simple sufficient case is coordinate selection: if Σ is stable and internally positive and $P = e^n_{\mathcal{I}}$, then $Q = e^n_{\mathcal{I}^c}$ (up to permutation), A_{PP} is a Hurwitz principal submatrix of the Hurwitz Metzler matrix A, and all factor realizations are internally positive. Equation (20) then applies.

The exact dc formula also yields an a priori upper bound under a strict dissipativity assumption. The setting is deliberately narrower than Theorem 2: the projection must be interface invariant, $-A$ must be uniformly coercive, and positive dominance is needed to promote the dc estimate to a uniform $\mathcal{H}_\infty$ certificate.

To state the result compactly, let $H := -A$, complete P to an orthogonal matrix $[P\ Q]$, and set

$$H_{PP} := P^THP, \qquad H_{PQ} := P^THQ,$$
$$H_{QP} := Q^THP, \qquad H_{QQ} := Q^THQ.$$

Define the Schur complement $S_Q := H_{QQ} - H_{QP}H_{PP}^{-1}H_{PQ}$ and the coupling factors

$$L_P := CPH_{PP}^{-1}H_{PQ}, \qquad R_P := H_{QP}H_{PP}^{-1}P^TB.$$

Theorem 4 (Coercive error certificate). *Let $\Sigma = (A, B, C, D)$ and let P define an IIPROM. Suppose*

$$\tfrac{1}{2}(H + H^T) \succeq \eta I_{n_x} \tag{21}$$

for some $\eta > 0$. Then H_{PP} and S_Q are nonsingular, and

$$\hat{\Sigma}_e(0) = L_PS_Q^{-1}R_P, \tag{22}$$

$$\|\hat{\Sigma}_e(0)\|_2 \leq \frac{\|L_P\|_2\|R_P\|_2}{\eta} \leq \frac{\|C\|_2\|B\|_2\|H_{PQ}\|_2\|H_{QP}\|_2}{\eta^3}. \tag{23}$$

If in addition Σ_e is positively dominated—in particular if the three factors in Theorem 3 are—then

$$\|\Sigma_e\|_\infty = \|\hat{\Sigma}_e(0)\|_2 \leq \frac{\|L_P\|_2\|R_P\|_2}{\eta}. \tag{24}$$

Proof. Interface invariance gives $Q^TB = 0$ and $CQ = 0$. Applying the block-inverse formula to $[P\ Q]^TH[P\ Q]$ and subtracting the reduced dc map $CPH_{PP}^{-1}P^TB$ from the full dc map $CH^{-1}B$ gives (22). Condition (21) implies $(H_{PP} + H_{PP}^T)/2 \succeq \eta I_r$, and hence $\|H_{PP}^{-1}\|_2 \leq 1/\eta$. For any y, choose $x = -H_{PP}^{-1}H_{PQ}y$. Then

$$\operatorname{Re}(y^*S_Qy) = \operatorname{Re}\left(\begin{bmatrix}x\\y\end{bmatrix}^*[P\ Q]^TH[P\ Q]\begin{bmatrix}x\\y\end{bmatrix}\right) \geq \eta\|y\|_2^2.$$

Thus $\|S_Q^{-1}\|_2 \leq 1/\eta$. Submultiplicativity gives the first inequality in (23); the second follows from $\|P\|_2 = \|Q\|_2 = 1$ and the bound on H_{PP}^{-1}. Positive dominance gives (24) by Proposition 5. □

Remark 4 (Collocated symmetric case). *If additionally $H = H^T$ and $C = B^T$, set $Z_P := H_{QP}H_{PP}^{-1}P^TB = L_P^T$. Then*

$S_Q \succ 0$, $\hat{\Sigma}_e(0) = Z_P^T S_Q^{-1} Z_P \succeq 0$, and under positive dominance $\|\Sigma_e\|_\infty \leq \|Z_P\|_2^2/\eta$. Indeed, symmetry gives $H_{PQ} = H_{QP}^T$ and $H_{PP} = H_{PP}^T$, so $L_P = CPH_{PP}^{-1}H_{PQ} = (H_{QP}H_{PP}^{-1}P^TB)^T = Z_P^T$ and $R_P = Z_P$; S_Q is a Schur complement of the positive-definite matrix $[P\ Q]^T H[P\ Q]$ and is therefore positive definite.

The word "a priori" here is projection dependent. Once a candidate P is specified, (24) can be evaluated without forming the full error system, computing S_Q^{-1}, or performing a frequency sweep. Partition-dependent a priori bounds of a similar kind are derived for clustering of leader–follower multi-agent systems in [12]. The bound (24) does not depend only on the target order, and it is not an approximation ratio relative to the globally optimal projection. Coercivity alone bounds only the dc error. Without positive dominance, an off-zero frequency peak can exceed (23). The coordinate-selection specialization of a stable internally positive system satisfies the needed positive-dominance condition, whereas a general dense projection or graph aggregation need not.

V. Model Reduction of MAS Using Graph Contractions

The preceding sections characterize admissible projections and their error, but do not prescribe how to choose the free subspace in Theorem 1. For multi-agent systems it is natural to require every reduced coordinate to represent one original agent or a connected aggregate, with actuated and measured agents retaining individual identities. Vertex partitions provide exactly this structure: their normalized characteristic matrices define orthogonal projections whose quotient graph gives a physical interpretation to the reduced dynamics. Requiring every interface vertex to be a singleton cell enforces the protected-subspace condition of Theorem 1, replacing the continuous search over projection subspaces with a combinatorial search over interpretable quotient graphs. When the positively dominated specialization of Section IV applies, the true $\mathcal{H}_\infty$ error equals the dc gain, so candidate contractions can be scored without a frequency sweep. This construction is a form of graph coarsening [21] and is related to clustering reductions [5]–[7], in particular to the positivity-preserving clustering of positive directed networks [10] and to dissimilarity-based clustering [8]; its distinctive constraint is exact preservation of every interface vertex.

A. Edge-Matching Graph Partitions

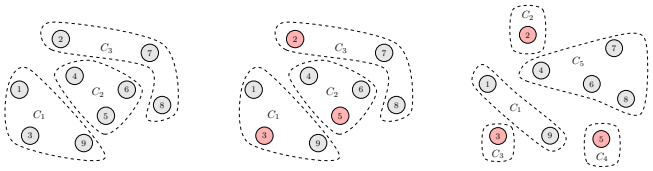
An *r-partition*, denoted as π_r, of a vertex set $\mathcal{V}$ is the partition $\{C_i\}_{i=1}^r$ of $\mathcal{V}$ into r cells such that

$$\bigcup_{i=1}^{r} C_i = \mathcal{V} \text{ and } C_i \cap C_j = \emptyset.$$

We call an r-partition *$\mathcal{I}$-invariant* for some $\mathcal{I} \subset \mathcal{V}$ if for each $i = 1, \ldots, r$, one has $|C_i \cap \mathcal{I}| \leq 1$. It is *strongly* $\mathcal{I}$-invariant if

$$C_i \cap \mathcal{I} \neq \emptyset \quad \Longrightarrow \quad |C_i| = 1. \tag{25}$$

Thus every interface vertex forms a singleton cell. We denote the set of strongly $\mathcal{I}$-invariant r-partitions by $S_r(\mathcal{V})$. Figure 4 illustrates the definitions.



(a) (b) (c)

Fig. 4. A set of $|\mathcal{V}| = 9$ nodes showing (a) a 3-partition, (b) an $\mathcal{I}$-invariant 3-partition with $\mathcal{I} = \{2, 3, 5\}$ (in red), and (c) a strongly $\mathcal{I}$-invariant 5-partition.

Given a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ and a subset of its edges $\mathcal{E}_s \subseteq \mathcal{E}$, the *edge-induced partition* $\pi(\mathcal{E}_s)$ is the collection of connected components of $(\mathcal{V}, \mathcal{E}_s)$ [22]. This is illustrated in Figure 5.

Given an r-partition $\pi = \{C_i\}_{i=1}^r$, we define $M(\pi) \in \mathbb{R}^{n \times r}$ as the *partition characteristic matrix* (PCM). Its entry $[M(\pi)]_{ij}$ equals 1 when $i \in C_j$ and 0 otherwise. The corresponding partition projection matrix (PPM) is $P(\pi) = M(\pi)(M(\pi)^T M(\pi))^{-1/2}$. For example, for the 3-partition in Figure 4a, we have

$$M(\pi) = \begin{bmatrix} 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 1 & 1 & 0 \end{bmatrix}^T.$$

For a MAS $\Sigma(\mathcal{G}, \mathcal{U}, \mathcal{Y})$ let $\mathcal{I} = \mathcal{U} \cup \mathcal{Y}$. If B and C select interface coordinates and $\pi \in S_r(\mathcal{V})$, then $P(\pi)P(\pi)^T$ fixes each interface basis vector. Consequently $B = P(\pi)P(\pi)^T B$ and $C = CP(\pi)P(\pi)^T$, so the partition PROM is interface invariant. Strong rather than ordinary invariance is required. Merging an interface vertex with a non-interface vertex changes its coordinate vector.

B. Stability of Positive MAS Contractions

For Metzler network dynamics, graph contractions admit a stability certificate that is distinct from the Euclidean dissipativity condition of Proposition 1. Recall that a Metzler matrix A is Hurwitz if and only if there exists a vector $v > 0$ such that $Av < 0$. A partition projection preserves this certificate whenever the positive vector is constant on every partition cell.

Proposition 6 (Positive-vector stability of MAS contractions)**.** *Let $A \in \mathbb{R}^{n \times n}$ be Metzler, and suppose there exists $v > 0$ such that $Av < 0$. Let $\pi = \{C_i\}_{i=1}^r$ be a partition with normalized characteristic projection $P = P(\pi) = M(\pi)(M(\pi)^T M(\pi))^{-1/2}$. If $v \in \operatorname{im} P$, equivalently, v is constant on every cell C_i, then the contracted matrix $A_r = P^T AP$ is Metzler and Hurwitz. More precisely, writing $v = Pv_r$, one has $v_r > 0$ and $A_r v_r = P^T Av < 0$.*

Proof. Because P is a normalized characteristic matrix, $P \geq 0$ and its columns have disjoint supports. The assumption $v \in \operatorname{im} P$ and $v > 0$ therefore imply that the unique vector $v_r = P^T v$ satisfying $v = Pv_r$ is strictly positive. For $i \neq j$, $[A_r]_{ij} = p_i^T A p_j \geq 0$, since $p_i, p_j \geq 0$ have disjoint supports and all entries of A connecting distinct coordinates are nonnegative; hence A_r is Metzler. Moreover, $A_r v_r = P^T APv_r = P^T Av < 0$, because $Av < 0$ and every

row of P^T is nonnegative and nonzero. The Metzler stability criterion then implies that A_r is Hurwitz. □

Remark 5 (Grounded-Laplacian networks)**.** *For $A = -(L + gI)$ with $g > 0$ and a graph Laplacian satisfying $L\mathbf{1} = 0$, choose $v = \mathbf{1}$. Then $Av = -g\mathbf{1} < 0$, and $\mathbf{1}$ is constant on every partition cell, so Proposition 6 applies to every partition. If $n_i = |C_i|$ and $v_r = [\sqrt{n_1}, \dots, \sqrt{n_r}]^T$, then $Pv_r = \mathbf{1}$ and $A_r v_r = -gv_r$. Thus every graph contraction is Hurwitz. Since A_r is Metzler and has the strictly positive eigenvector v_r, its spectral abscissa is $-g$. Hence the contraction preserves the decay margin exactly.*

This criterion is complementary to Proposition 1: it is specific to nonnegative partition projections and Metzler dynamics, but does not require Euclidean dissipativity and therefore also applies to nonsymmetric positive network models. Positive dominance itself remains an input–output property of an already stable transfer matrix; Proposition 6 supplies the internal stability certificate needed for this class of MAS contractions.

The natural goal is to find the partition $\pi^* \in S_r(\mathcal{V})$ that minimizes $\|\Sigma - \Sigma_r(P(\pi))\|_\infty$. Finding such an optimal partition, however, is numerically intractable for a moderate number of nodes: the number of r-partitions is the Stirling number of the second kind, which for $r \ll n$ is asymptotically $S(n, r) \sim \frac{r^n}{r!}$ [23, p.18].

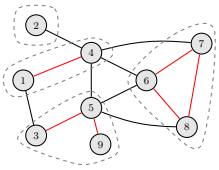


Fig. 5. An edge-induced partition of a connected graph $\mathcal{G}$. The edges in $\mathcal{E}_s$ are drawn in red and the resulting connected components of the induced graph $\hat{\mathcal{G}} = (\mathcal{V}, \mathcal{E}_s)$ are outlined with dashed lines. Thus the edge-induced partition $\pi(\mathcal{E}_s)$ yields a 4-partition.

C. Sequential Greedy Coarsening

Since exact minimization over $S_r(\mathcal{V})$ is combinatorially intractable, we adopt a greedy heuristic that builds the target partition one vertex-merge at a time. A naive approach would select the $n - r$ best edges simultaneously, but this can create cycles and need not produce exactly r cells; it also ignores interactions between successive merges. Algorithm 1 below addresses both issues: it performs one admissible merge at a time, recomputing every candidate score after each step. The default score is the true error $J(P) = \|\Sigma - \Sigma_r(P)\|_\infty$, directly addressing cases in which the product bound is conservative. When exact norm evaluation is too costly, b_{prod} in (16) may be substituted, but only after the stability tests required by Corollary 1.

The procedure is formalized in Algorithm 1. Starting from the finest partition (each vertex its own cell), at each iteration all edge-adjacent cell pairs are enumerated; merges that would violate the strong interface-invariance constraint (25) or yield an unstable reduced model are discarded, and the remaining candidates are scored. The algorithm commits to the merge with the best score and repeats until the target order r is reached or no admissible merge remains.

Algorithm 1 Sequential interface-preserving edge coarsening

Input: MAS $\Sigma(\mathcal{G}, \mathcal{U}, \mathcal{Y})$ of order n, candidate edges $\mathcal{E}_c \subseteq \mathcal{E}$, target order $r \geq |\mathcal{I}|$, score $J(\pi') := \|\Sigma - \Sigma_{|\pi'|}(P(\pi'))\|_\infty$ or $b_{\text{prod}}(\Sigma, P(\pi'))$, optional Lyapunov X.

1) **Initialize.** Set $\pi \leftarrow \{\{v\} : v \in \mathcal{V}\}$. Test (7); record μ if satisfied, else validate X. For Metzler dynamics, check whether $v > 0$ with $Av < 0$ is available (Prop. 6).
2) **Repeat while** $|\pi| > r$**:**
   (a) *Generate.* Form all π' by merging adjacent cells via edges in $\mathcal{E}_c$.
   (b) *Filter.* Discard π' violating (25). Certify stability via Prop. 1; for Metzler dynamics use Prop. 6 when $v \in \operatorname{im} P(\pi')$. Remaining uncertified candidates undergo a direct Hurwitz test; unstable ones are discarded. If $J = b_{\text{prod}}$, also discard candidates with unstable factors.
   (c) *Select.* If no admissible π' remains, return `early-stop`. Otherwise set $\pi \leftarrow \operatorname{argmin}_{\pi'} J(\pi')$ and record $J(\pi)$, b_{prod}, and the certificate.

Output: $\Sigma_{|\pi|}(P(\pi))$, termination status, and per-step scores and certificates.

The stability certificates used in Algorithm 1 are Proposition 1 and, for positive network contractions, the certificate of Proposition 6 tailored to nonnegative partition projections.

Taking $\mathcal{E}_c = \mathcal{E}$ gives the richest local search and is the default in the supplied implementation. The merge objective is not known to be submodular, so data-independent greedy approximation ratios do not apply. The following result instead provides a computable a posteriori certificate for every run that reaches the requested order.

Proposition 7 (A posteriori greedy certificate)**.** *Let $\sigma^{\mathrm{H}}_{r+1}(\Sigma)$ be the $(r+1)$st Hankel singular value of the stable full-order system, with value zero when its McMillan degree is at most r. Suppose Algorithm 1 reaches the requested order with a stable sequence $\{\Sigma^{\mathrm{g}}_{(k)}\}_{k=0}^{n-r}$, and define*

$$J^{\mathrm{g}}_r := \|\Sigma - \Sigma^{\mathrm{g}}_{(n-r)}\|_\infty, \qquad U^{\mathrm{g}}_r := \sum_{k=1}^{n-r} \|\Sigma^{\mathrm{g}}_{(k)} - \Sigma^{\mathrm{g}}_{(k-1)}\|_\infty.$$

Let $J^\star_r$ be the optimal value of Problem 1, and let $V^\star_r$ be the optimal value of Problem 2. Then

$$\sigma^{\mathrm{H}}_{r+1}(\Sigma) \leq J^\star_r \leq J^{\mathrm{g}}_r \leq U^{\mathrm{g}}_r, \qquad \sigma^{\mathrm{H}}_{r+1}(\Sigma) \leq V^\star_r \leq U^{\mathrm{g}}_r. \quad (26)$$

Consequently, if $\sigma^{\mathrm{H}}_{r+1}(\Sigma) > 0$, the returned model satisfies the certified approximation factors

$$\frac{J^{\mathrm{g}}_r}{J^\star_r} \leq \frac{J^{\mathrm{g}}_r}{\sigma^{\mathrm{H}}_{r+1}(\Sigma)}, \qquad \frac{U^{\mathrm{g}}_r}{V^\star_r} \leq \frac{U^{\mathrm{g}}_r}{\sigma^{\mathrm{H}}_{r+1}(\Sigma)}. \quad (27)$$

Proof. For every stable transfer matrix $\widetilde{\Sigma}_r$ of McMillan degree at most r, the Hankel-operator approximation theorem gives

$$\|\Sigma - \widetilde{\Sigma}_r\|_\infty \geq \|\mathcal{H}_\Sigma - \mathcal{H}_{\widetilde{\Sigma}_r}\| \geq \sigma_{r+1}^{\mathrm{H}}(\Sigma)$$

[3], [24]. This applies to every feasible final model in both problems. The triangle inequality gives $J_r^{\mathrm{g}} \leq U_r^{\mathrm{g}}$ and, for every feasible sequence, its accumulated cost is no smaller than its final error. Finally, the greedy output is feasible for both problems whenever it terminates. These observations give (26); division by a positive lower bound gives (27). □

Proposition 7 is an a posteriori result; it does not establish a uniform approximation ratio. Algorithm 1 returns the certified upper bound U_r^{g}, while the Hankel singular value supplies a universal lower bound on either global optimum. The lower bound is taken over all stable order-r approximants, including dense unstructured ones, and can be much smaller than the best interface-preserving error when the plant is nearly reducible in unstructured coordinates—as in the thermal network of Section VI, where the interface-preserving error exceeds $\sigma_{r+1}^{\mathrm{H}}$ by a factor of at least 277 already at $r = 60$. The certificate is therefore a worst-case safety net rather than a typical performance prediction.

VI. Case Study: A 65,566-State Chip Thermal Network

Dynamic thermal management of a processor acts on the power of individual functional units and reads their temperatures, so a compact thermal model is useful only if every unit keeps its own input and output. The network below is internally positive, symmetric, and strictly dissipative, so every result of Sections II–V applies to it. The structured baseline is Kron reduction of the same interface, known in thermal modeling as static condensation or Guyan reduction [25]: it keeps the interface nodes and $r - |\mathcal{I}|$ interior nodes and eliminates the rest through a Schur complement. It is a Petrov–Galerkin PROM that preserves the interface and matches the dc gain exactly, but it is not an orthogonal projection, so Sections III–IV do not apply to it. We also compare against coordinate deletion and balanced truncation (BT), and report the true relative error $\|\Sigma - \Sigma_r\|_\infty / \|\Sigma\|_\infty$.

A. Model and Reduction Methods

The network follows the HotSpot methodology [26], but not its solver, for the public Alpha EV6 floorplan and GCC power trace. A $150\,\mu$m silicon die ($k = 130\,\mathrm{W/(m\,K)}$, $1.63\,\mathrm{MJ/(m^3\,K)}$) is discretized into a 256×256 grid. Each of the 30 functional units has one active-layer node ($10\,\mu$m) coupled to the cells beneath it through a $20\,\mu$m interface layer ($4\,\mathrm{W/(m\,K)}$), and the grid is cooled through a uniform $0.1\,\mathrm{K/W}$ package resistance, giving $n = 65{,}566$. With conductance matrix G (a grounded Laplacian), diagonal heat capacities C_{th}, unit-node selector S, and energy coordinates $x = C_{\mathrm{th}}^{1/2}(T - T_{\mathrm{amb}})$, the model is

$$\dot{x} = -Hx + Bu, \qquad y = B^T x, \tag{28}$$

with $H = C_{\mathrm{th}}^{-1/2} G C_{\mathrm{th}}^{-1/2}$ and $B = C_{\mathrm{th}}^{-1/2} S$; the inputs are the unit powers and the outputs are the unit temperature rises.

TABLE I
EV6 thermal network: accuracy, structure, and cost

r	Method	Err. (%)	ω_{pk}	Edges	Pos.	Int.	Cap.	Time
402	IIPROM	24.0	0	1,199	✓	✓	100%	0.06
	Kron-dc	34.0	466	12,144	✓	✓	6.8%	1.2
	Kron-rep	42.3	633	51,644	✓	✓	6.8%	1.3
	Deletion	58.5	0	372	✓	✓	6.8%	< 0.01
	BT	10^{-10}	0.6	80,601	×	×	–	20.6
6402	IIPROM	1.38	0	19,702	✓	✓	100%	0.15
	Kron-dc	3.10	406	31,902	✓	✓	15.4%	32.6
	Kron-rep	5.13	209	284,305	✓	✓	15.4%	22.7
	Deletion	47.0	0	14,892	✓	✓	15.4%	< 0.01

Pos.: A_r Metzler and $B_r \geq 0$. Int.: interface exactly preserved. Cap.: fraction of heat capacity retained. Time: seconds, excluding the 0.46 s dc factorization shared by IIPROM and Kron-dc.

Here $H \succeq \eta I$ with $\eta = 149.6\,\mathrm{s}^{-1}$, $\|\Sigma\|_\infty = \|S^T G^{-1} S\|_2 = 44.9\,\mathrm{K/W}$, and $r_{\min} = 30$ by Theorem 1.

IIPROM. Every unit node is a singleton cell, and $P(\pi) = C_{\mathrm{th}}^{1/2} M (M^T C_{\mathrm{th}} M)^{-1/2}$, $M = M(\pi)$, is the capacitance-weighted partition projection of Section V: it is orthonormal and nonnegative and contains $\mathcal{S}_{\mathrm{io}}$. In temperature coordinates the reduced model is an RC network on the quotient graph whose states are cluster-mean temperatures. The partition is built by a one-pass divisive heuristic rather than by Algorithm 1. Starting from one cluster per unit, it repeatedly splits into quadrants the cluster with the largest capacitance-weighted spread $\sum_{i \in C} c_i \|T_i - \bar{T}_C\|_2^2$ of the dc temperature field $T = G^{-1} S$; these spreads sum to $\|(I - PP^T) X_0\|_F^2$, $X_0 = H^{-1} B$, the dc response discarded by aggregation. Cells never cross unit boundaries, so the smallest order is $r = 60$.

Baselines. Kron reduction retains either the $r - 30$ cells of highest dc-influence $\|T_i\|_2$ (Kron-dc) or one representative cell per IIPROM cluster (Kron-rep). Coordinate deletion keeps the Kron-rep nodes but truncates orthogonally, $P = e_{\mathcal{I}}$, the setting of Theorem 3. Since H is symmetric and $C = B^T$, BT is the orthogonal PROM onto the r leading Gramian eigenvectors, whose range does not contain $\mathcal{S}_{\mathrm{io}}$. The Gramian factor is computed by low-rank ADI [27] (rank 1440, 20.6 s), so BT is reported for $r \leq 801$. All reduced models are stable with decay margin at least η, the orthogonal ones by Proposition 1(i) and Kron because a Schur complement of H satisfies the same bound. Errors are computed from sparse factorizations of the full model at 90 log-spaced frequencies and at $\omega = 0$, refining every off-dc peak.

B. Results

Structure. The IIPROM is a sparse lumped RC network (Fig. 6, Table I). It has about three edges per state at every order and retains all of the heat capacity, and each of its states is the mean temperature of an identifiable region of the die. The Kron-reduced model is also an RC network on physical nodes, because a Schur complement of an M-matrix is an M-matrix. However, it discards 85–94% of the heat capacity and fills in: at $r = 402$, Kron-dc has 10 times and Kron-rep 43 times as many edges as the IIPROM. BT yields a complete graph with a non-Metzler state matrix and a mixed-sign input matrix. Its range only approximates the interface ($\|(I - PP^T)B\|_2 / \|B\|_2 = 4.7\%$ at $r = 60$), and its coordinates are global modes (Fig. 6d).

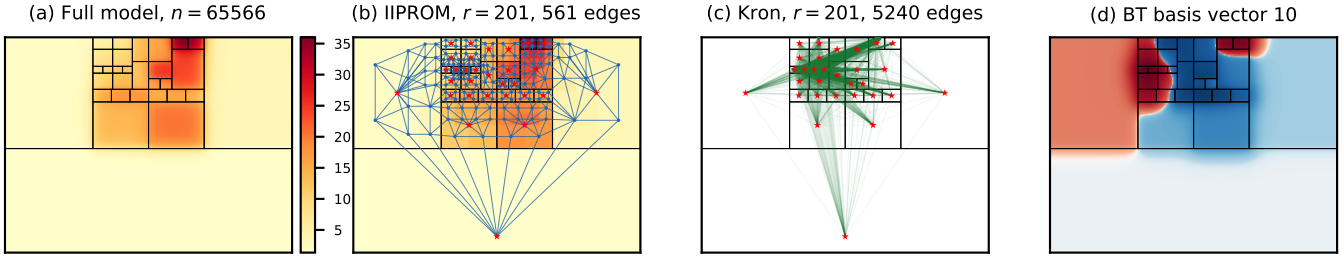


Fig. 6. EV6 thermal network and reduced models at $r = 201$, with the uniform lower part of the L2 cache cropped; red stars mark the 30 unit (interface) nodes. (a) Steady-state temperature rise (K) under the mean GCC power (40.2 W). (b) IIPROM: reduced graph (blue) over the cluster-mean temperatures. (c) Kron-dc: retained cells and their 5,240-edge reduced graph. (d) A BT basis vector (symmetric log scale), which corresponds to no physical location.

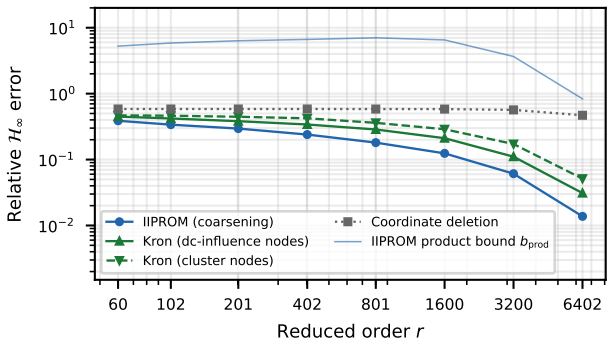


Fig. 7. EV6 thermal network: relative $\mathcal{H}_\infty$ error of the structure-preserving reductions vs. reduced order, with the IIPROM product bound of Corollary 1.

Accuracy. Among the structure-preserving reductions (Fig. 7), the IIPROM is the most accurate at every order from $r = 60$ to $r = 6{,}402$. The better Kron variant has 1.16 times the IIPROM error at $r = 60$ (44.9% vs. 38.8%), 1.6 times at $r = 801$, and 2.25 times at $r = 6{,}402$ (3.10% vs. 1.38%). Coordinate deletion stays at 47–58%, because deleting a cell grounds it to ambient. Per edge the advantage widens: the IIPROM reaches 18% error with 2,466 edges and 6.1% with 9,981, whereas Kron-dc needs 22,036 and 35,913 edges for 28.6% and 11.1%. BT is far more accurate, with 0.29% at $r = 60$ and machine precision from $r = 402$, and it tracks the Hankel lower bound within a factor of about 2.5. The full model is therefore almost exactly reducible in unstructured coordinates, and the gap between BT and every node-based model is the price of keeping physical structure. For the same reason, the a posteriori factor of Proposition 7 is uninformative here: $J_r/\sigma^{\mathrm{H}}_{r+1} \geq 277$ already at $r = 60$.

Where the error lives. The IIPROM and deletion errors peak at $\omega = 0$ at every tested order (Table I). For deletion this is guaranteed by Theorem 3. For the IIPROM it is an observation rather than a certified property, since Q has mixed signs, but it means that the dc-gain error, obtained from a single sparse solve, equals the $\mathcal{H}_\infty$ error. Kron reduction is exact at dc, and its error peaks at 209–812 rad/s, above η, where the discarded heat capacity matters. Kron therefore gives exact steady-state temperatures for power profiles slower than about $1/\eta = 6.7$ ms, whereas the IIPROM has the smaller worst-case error over all frequencies.

Certificates and cost. The product bound b_{prod} holds at every order, but it is conservative by a factor of 14 at $r = 60$ and 60 at $r = 6{,}402$, because it ignores how the input and output couplings align. The coercive bound of Remark 4 exceeds $\|\Sigma\|_\infty$ and is uninformative here. Once the dc field is available (0.46 s), an IIPROM is built in 0.02–0.15 s by aggregation alone. Kron reduction takes 0.3 s at $r = 60$ and 33 s at $r = 6{,}402$, and BT needs the Gramian and yields dense models. The complete study runs in 16 minutes on a laptop.

VII. Conclusions

We characterized interface-invariant orthogonal projections as those whose retained subspace contains $\mathrm{im}[B\ C^T]$, and factored the reduction error into input-coupling, internal-dynamics, and output-coupling systems. The factorization gives a submultiplicative $\mathcal{H}_\infty$ bound that becomes an exact dc-gain expression for positively dominated factors; strict dissipativity certifies stability, and a Hankel singular value gives an a posteriori approximation factor.

On a 65,566-state processor thermal network with 30 functional-unit interfaces, interface-preserving graph coarsening produced sparse, positive RC models that retain all of the heat capacity. At every tested order they were more accurate in $\mathcal{H}_\infty$ than Kron reduction of the same interface, by up to a factor of 2.25 and with up to 43 times fewer edges. Balanced truncation was far more accurate but produced dense models whose coordinates mix the interface; this gap is the price of keeping physical structure.

Appendix A
Proof of Theorem 2

Proof. The proof has two steps: an expression for $\hat{\Sigma}_e$ in terms of the discarded directions Q, and its factorization. Both use the block form of Φ in the coordinates $U = [P\ \ Q]$,

$$N = U^T \Phi U = \begin{bmatrix} sI_r - A_{PP} & -A_{PQ} \\ -A_{QP} & sI - A_{QQ} \end{bmatrix},$$

with blocks N_{ij}. Since U is orthogonal, $N^{-1} = U^T \Phi^{-1} U$; in particular, $(N^{-1})_{22} = Q^T \Phi^{-1} Q$.

Error transfer matrix. The feedthrough D cancels, and $\hat{\Sigma}_r = CP\,N_{11}^{-1} P^T B$, so $\hat{\Sigma}_e = C(\Phi^{-1} - P N_{11}^{-1} P^T)B$. With $S = N_{22} - N_{21} N_{11}^{-1} N_{12}$, the partitioned inverse [28] gives $N^{-1} - \mathrm{diag}(N_{11}^{-1}, 0) = [-N_{11}^{-1} N_{12};\ I]\, S^{-1}\, [-N_{21} N_{11}^{-1}\ \ I]$. The two

outer factors times S^{-1} are the second block column and row of N^{-1}, and $(N^{-1})_{22} = S^{-1}$, so the right-hand side equals $(N^{-1})_{:,2}\,(N^{-1})_{22}^{-1}\,(N^{-1})_{2,:}$. Multiplying by U on the left and U^T on the right, and using $UN^{-1} = \Phi^{-1}U$, gives

$$\Phi^{-1} - PN_{11}^{-1}P^T = \Phi^{-1}Q(Q^T\Phi^{-1}Q)^{-1}Q^T\Phi^{-1},$$

which for symmetric positive-definite Φ is Khatri's lemma [29]. Substituting into $\hat{\Sigma}_e$ proves (13).

Factorization. Let $\hat{\Delta} = Q^T\Phi^{-1}Q$, $\hat{\Theta} = C\Phi^{-1}Q\hat{\Delta}^{-1}$, and $\hat{\Gamma} = \hat{\Delta}^{-1}Q^T\Phi^{-1}B$, so that $\hat{\Sigma}_e = \hat{\Theta}\hat{\Delta}\hat{\Gamma}$ by (13); $\hat{\Delta}$ has the realization $(A, Q, Q^T, 0)$ directly. The off-diagonal blocks of $NN^{-1} = I$ and $N^{-1}N = I$ give $(N^{-1})_{12} = -N_{11}^{-1}N_{12}(N^{-1})_{22}$ and $(N^{-1})_{21} = -(N^{-1})_{22}N_{21}N_{11}^{-1}$. Inserting $I = PP^T + QQ^T$ next to Φ^{-1}, and using $P^T\Phi^{-1}Q = (N^{-1})_{12}$ and $Q^T\Phi^{-1}P = (N^{-1})_{21}$, yields

$$\Phi^{-1}Q\hat{\Delta}^{-1} = P(sI_r - A_{PP})^{-1}A_{PQ} + Q,$$
$$\hat{\Delta}^{-1}Q^T\Phi^{-1} = A_{QP}(sI_r - A_{PP})^{-1}P^T + Q^T.$$

Hence $\hat{\Theta} = CP(sI_r - A_{PP})^{-1}A_{PQ} + CQ$ and $\hat{\Gamma} = A_{QP}(sI_r - A_{PP})^{-1}P^TB + Q^TB$, the transfer matrices of Θ and Γ in (15). □

ACKNOWLEDGMENT

During the preparation of this work, the authors used Anthropic Claude to structure some of the proofs, to help develop the numerical case study, and to assist with editing the manuscript. After using this tool/service, the authors reviewed, restructured, and edited the content as needed and take full responsibility for the content of the publication.

REFERENCES

[1] G. E. Dullerud and F. Paganini, *A Course in Robust Control Theory: A Convex Approach*. Springer, 2000.
[2] B. C. Moore, "Principal component analysis in linear systems: Controllability, observability, and model reduction," *IEEE Trans. Autom. Control*, vol. 26, no. 1, pp. 17–32, 1981.
[3] A. C. Antoulas, *Approximation of Large-Scale Dynamical Systems*. SIAM, 2005.
[4] V. Simoncini, "Computational methods for linear matrix equations," *SIAM Rev.*, vol. 58, no. 3, pp. 377–441, 2016.
[5] T. Ishizaki, K. Kashima, J.-i. Imura, and K. Aihara, "Model reduction and clusterization of large-scale bidirectional networks," *IEEE Trans. Autom. Control*, vol. 59, no. 1, pp. 48–63, 2014.
[6] N. Monshizadeh, H. L. Trentelman, and M. K. Camlibel, "Projection-based model reduction of multi-agent systems using graph partitions," *IEEE Trans. Control Netw. Syst.*, vol. 1, no. 2, pp. 145–154, 2014.
[7] B. Besselink, H. Sandberg, and K. H. Johansson, "Clustering-based model reduction of networked passive systems," *IEEE Trans. Autom. Control*, vol. 61, no. 10, pp. 2958–2973, 2016.
[8] X. Cheng, Y. Kawano, and J. M. A. Scherpen, "Model reduction of multiagent systems using dissimilarity-based clustering," *IEEE Trans. Autom. Control*, vol. 64, no. 4, pp. 1663–1670, 2019.
[9] X. Cheng and J. M. A. Scherpen, "Model reduction methods for complex network systems," *Annu. Rev. Control Robot. Auton. Syst.*, vol. 4, pp. 425–453, 2021.
[10] T. Ishizaki, K. Kashima, A. Girard, J.-i. Imura, L. Chen, and K. Aihara, "Clustered model reduction of positive directed networks," *Automatica*, vol. 59, pp. 238–247, 2015.
[11] N. Leiter and D. Zelazo, "Graph-based model reduction of the controlled consensus protocol," in *Proc. 20th IFAC World Congr.*, 2017, pp. 9456–9461.
[12] H.-J. Jongsma, P. Mlinarić, S. Grundel, P. Benner, and H. L. Trentelman, "Model reduction of linear multi-agent systems by clustering with H_2/H_∞ error bounds," *Math. Control Signals Syst.*, vol. 30, no. 1, 2018, art. no. 6.
[13] X. Cheng, J. M. A. Scherpen, and B. Besselink, "Balanced truncation of networked linear passive systems," *Automatica*, vol. 104, pp. 17–25, 2019.
[14] R. Sabbagh and D. Abou Jaoude, "Model reduction of consensus network systems via selection of optimal edge weights and nodal time-scales," in *Proc. Amer. Control Conf. (ACC)*, 2022, pp. 1859–1866.
[15] N. van de Wouw, B. Besselink, M. F. Shakib, L. Janssen, L. Poort, G. Scarciotti, and R. H. B. Fey, "Structure-preserving model reduction of interconnected dynamical systems," *Annu. Rev. Control*, vol. 60, 2025, art. no. 101016.
[16] X. Cheng, "An optimal projection framework for structure-preserving model reduction of linear systems," *IEEE Trans. Autom. Control*, 2026, early access, doi: 10.1109/TAC.2026.3747707.
[17] F. Dörfler and F. Bullo, "Kron reduction of graphs with applications to electrical networks," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 60, no. 1, pp. 150–163, 2013.
[18] A. Rantzer, "Optimizing positively dominated systems," in *Proc. 51st IEEE Conf. Decis. Control (CDC)*, 2012, pp. 272–277.
[19] T. Tanaka and C. Langbort, "The bounded real lemma for internally positive systems and H-infinity structured static state feedback," *IEEE Trans. Autom. Control*, vol. 56, no. 9, pp. 2218–2223, 2011.
[20] L. Farina and S. Rinaldi, *Positive Linear Systems: Theory and Applications*. Wiley, 2000.
[21] A. Loukas, "Graph reduction with spectral and cut guarantees," *J. Mach. Learn. Res.*, vol. 20, no. 116, pp. 1–42, 2019.
[22] N. Leiter and D. Zelazo, "Edge-matching graph contractions and their interlacing properties," *Linear Algebra Appl.*, vol. 612, pp. 289–317, 2021.
[23] H. S. Wilf, *generatingfunctionology*, 2nd ed. Academic Press, 1994.
[24] K. Glover, "All optimal Hankel-norm approximations of linear multivariable systems and their L^∞-error bounds," *Int. J. Control*, vol. 39, no. 6, pp. 1115–1193, 1984.
[25] R. J. Guyan, "Reduction of stiffness and mass matrices," *AIAA J.*, vol. 3, no. 2, p. 380, 1965.
[26] W. Huang, S. Ghosh, S. Velusamy, K. Sankaranarayanan, K. Skadron, and M. R. Stan, "HotSpot: A compact thermal modeling methodology for early-stage VLSI design," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 14, no. 5, pp. 501–513, 2006.
[27] T. Penzl, "A cyclic low-rank Smith method for large sparse Lyapunov equations," *SIAM J. Sci. Comput.*, vol. 21, no. 4, pp. 1401–1418, 2000.
[28] D. S. Bernstein, *Matrix Mathematics: Theory, Facts, and Formulas*, 2nd ed. Princeton Univ. Press, 2009.
[29] C. G. Khatri, "A note on a MANOVA model applied to problems in growth curve," *Ann. Inst. Statist. Math.*, vol. 18, no. 1, pp. 75–86, 1966.