



MAX PLANCK INSTITUTE  
FOR DYNAMICS OF COMPLEX  
TECHNICAL SYSTEMS  
MAGDEBURG



COMPUTATIONAL METHODS IN  
SYSTEMS AND CONTROL THEORY

# Comprehensive Memory-Bound Simulations on Single Board Computers

Christian Himpe, Tobias Leibner\*, Stephan Rave\*

Max Planck Institute Magdeburg  
\* University of Münster

2<sup>nd</sup> Workshop on Power Aware Computing (#paco2017)  
2017-07-07



- Single board computers (SBC) are widely available!
- Increasing performance of SBC (mobile) CPUs!
- SBCs have vector units for number crunching!





1. (Nonlinear) Model Order Reduction
2. Memory Economic Computation (HAPOD)
3. Numerical Experiments (with SBCs)

- Differential Equation System Models
- High-Dimensional State Space
- Fast(er) Solves
- Repeated Solution (Inputs, Parameters)

Special Case: (My) Nonlinear MOR

- Data-Driven
- Empirical Gramians
- No Hyperreduction
- Application: Gas network simulation (MathEnergy)

Nonlinear Ordinary Differential Equation (ODE):

$$\dot{x}(t) = f(t, x(t))$$

Nonlinear Input-Output System:

$$\dot{x}(t) = f(x(t), u(t))$$

$$y(t) = g(x(t), u(t))$$

Reduced Order Model (ROM):

$$\dot{x}_r(t) = f_r(x_r(t), u(t))$$

$$y_r(t) = g_r(x_r(t), u(t))$$

ROM Properties:

- $\dim(x_r(t)) \ll \dim(x(t))$
- $\|y - y_r\| \ll 1$

Projection-Based ROM:

$$\dot{x}_r(t) = V_1 f(U_1 x_r(t), u(t))$$

$$y_r(t) = g(U_1 x_r(t), u(t))$$

Truncated Projection Properties:

- $V_1 U_1 = \mathbb{1}$
- $\text{rank}(U_1) = \text{rank}(V_1) \ll \dim(x(t))$



Linear Input-Output System:

$$\dot{x}(t) = Ax(t) + Bu(t)$$

$$y(t) = Cx(t)$$

Linear ROM:

$$\dot{x}_r(t) = A_r x_r(t) + B_r u(t)$$

$$y_r(t) = C_r x_r(t)$$

Linear Projection-Based ROM:

$$\dot{x}_r(t) = (V_1 A U_1) x_r(t) + (V_1 B) u(t)$$

$$y_r(t) = (C U_1) x_r(t)$$





# Balanced Truncation<sup>1</sup>

Controllability Gramian Matrix & Observability Gramian Matrix:

$$W_C := \int_0^\infty e^{At} B B^T e^{A^T t} dt, \quad W_O := \int_0^\infty e^{A^T t} C^T C e^{At} dt$$

Balancing and Truncation<sup>2</sup>:

$$W_C = L_C L_C^T, \quad W_O = L_O L_O^T$$

$$L_O L_C^T \stackrel{\text{SVD}}{=} U D V$$

$$U_1 = U_{*,k}, \quad V_1 = V_{k,*}, \quad \varphi(D, k) < \varepsilon$$

---

<sup>1</sup>B.C. Moore. **Principal component analysis in linear systems: controllability, observability, and model reduction.** IEEE Transactions on Automatic Control 26(1): 17–32, 1981.

<sup>2</sup>M.S. Tombs and I. Postlethwaite. **Truncated balanced realization of stable, non-minimal state-space systems.** International Journal of Control 46: 1349–1330, 1989.



Cross Gramian Matrix:

$$W_X := \int_0^{\infty} e^{At} BC e^{At} dt$$

Approximate Balancing and Truncation<sup>4</sup>:

$$W_X \stackrel{\text{SVD}}{=} UDV$$
$$U_1 = U_{*,k}, \quad V_1 = U_1^T, \quad \tilde{\varphi}(D, k) < \varepsilon$$

- $\dim(u(t)) \stackrel{!}{=} \dim(y(t))$
- $C(1s - A)^{-1}B = (C(1s - A)^{-1}B)^T \Rightarrow W_X = BT$

---

<sup>3</sup>K.V. Fernando and H. Nicholson. **On the Structure of Balanced and Other Principal Representations of SISO Systems.** IEEE Transactions on Automatic Control 28(2): 228–231, 1983.

<sup>4</sup>D.C. Sorensen and A.C. Antoulas. **Projection methods for balanced model reduction.** Rice University TR01-03: 1–18, 2001.



Linear Cross Gramian<sup>5 6</sup>:

$$\begin{aligned}W_X &= \int_0^{\infty} e^{At} BC e^{At} dt \\&= \int_0^{\infty} (e^{At} B)(C e^{At}) dt \\&= \int_0^{\infty} (e^{At} B)(e^{A^T t} C^T)^T dt\end{aligned}$$

- Primal impulse response:  $g(t) = e^{At} B$
- Adjoint impulse response:  $\bar{g}(t) = e^{A^T t} C^T$

<sup>5</sup>K.V. Fernando and H. Nicholson. **On the Cross-Gramian for Symmetric MIMO Systems**. IEEE Transactions on Circuits and Systems 32(5): 487–489, 1985.

<sup>6</sup>H.R. Shaker. **Generalized Cross-Gramian for Linear Systems**. Proceedings of the IEEE Conference on Industrial Electronics and Applications: 749–751, 2012.



Empirical Cross Gramian Matrix:

$$\widehat{W}_X := \sum_{m=1}^{\dim(u(t))} \int_0^\infty \Psi^m(t) dt$$
$$\Psi_{ij}^m(t) = (x_i^m(t) - \bar{x}_i^m)(y_m^j(t) - \bar{y}_m^j)$$

- State trajectories ( $m$ -th perturbed input):  $x^m(t)$
- Output trajectories ( $j$ -th perturbed init state):  $y^j(t)$

---

<sup>7</sup>S. Streif, R. Findeisen and E. Bullinger. **Relating Cross Gramians and Sensitivity Analysis in Systems Biology**. Theory of Networks and Systems 10.4: 437–442, 2006.

<sup>8</sup>C. H. and M. Ohlberger. **Cross-Gramian Based Combined State and Parameter Reduction for Large-Scale Control Systems**. Mathematical Problems in Engineering 2014: 1–13, 2014.



Column-Wise Representation:

$$\widehat{W}_X = (w_{X,1} \quad w_{X,2} \quad \dots \quad w_{X,N})$$

Empirical Cross Gramian Column:

$$w_{X,j} = \sum_{m=1}^{\dim(u(t))} \int_0^{\infty} \psi^{jm}(t) dt$$

$$\psi_i^{jm}(t) = (x_i^m(t) - \bar{x}_i^m)(y_m^j(t) - \bar{y}_m^j)$$

The empirical cross Gramian is the only empirical Gramian that can be computed column-wise, because it is actually not a Gramian matrix!

---

<sup>9</sup>C. H. Combined State and Parameter Reduction for Nonlinear Systems with an Application in Neuroscience. Sierke Verlag Göttingen, 2017.

Aim: Truncated SVD of Empirical Cross Gramian:

$$W_X \stackrel{\text{SVD}}{=} UDV \rightarrow U_1$$

If only I could ...

- ... compute an SVD for a few columns at a time.
- ... enforce a projection error for  $U_1$ .
- ... guarantee correctness.



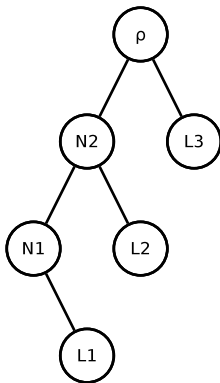
Hierarchical Approximate Proper Orthogonal Decomposition<sup>10</sup>(HAPOD):

$$\hat{u}_0 := \{\},$$
$$[\omega_s, \hat{u}_{s-1}] \stackrel{\text{SVD}}{=} u_s d_s v_s \rightarrow \hat{u}_s := u_s \hat{d}_s,$$
$$\hat{d}_{s,ii} = \begin{cases} d_{s,ii} & d_{s,ii} < \varepsilon^2 K_s \sqrt{\frac{\sum_{j=1}^s K_j}{S}} \\ 0 & \text{else} \end{cases}$$
$$U_1 := \hat{u}_S \rightarrow V_1 = U_1^T.$$

- Incremental HAPOD variant
- Proper Orthogonal Decomposition (POD)
  - Basically the argument's left singular values

---

<sup>10</sup>HAPOD Matlab implementation: <http://git.io/hapod>



- Compute as many columns as fit into memory as leafs  $L$ .
- Compute Sub-POD at nodes  $N$ .
- Go up the tree.
- The root POD  $\rho$  yields the HAPOD.





- Partitions distributed as leafs of some rooted tree.
- Natural per depth level parallelization.
- Projection-error-based, NOT rank-based.
- Rigorous and tight error and mode bounds.
- Only low-rank quantities are communicated.
- Compute POD, SVD or PCA.
- Use custom SVD implementation for “local” PODs.
- Relaxation parameter (accuracy vs speed)

---

<sup>11</sup>C. Himpe, T. Leibner and S. Rave. **Hierarchical Approximate Proper Orthogonal Decomposition**. arXiv e-prints math.NA: 1607.05210, 2017.



Hyperbolic Network Model (HNM)<sup>13</sup>:

$$\begin{aligned}\dot{x}(t) &= A \tanh(Kx(t)) + Bu(t) \\ y(t) &= Cx(t)\end{aligned}$$

- $\dim(x(t)) = 1024$
- $\dim(u(t)) = \dim(y(t)) = 1$

Numerical Experiment:

1. Compute empirical cross Gramian columns via `emgr`
2. Compute left singular vectors via `HAPOD`

---

<sup>13</sup>Y. Quan, H. Zhang and L. Cai. **Modeling and Control Based on a New Neural Network Model**. Proceedings of the American Control Conference 3: 1928–1929, 2001.



## Empirical Gramians:

- Empirical Controllability Gramian
- Empirical Observability Gramian
- Empirical Linear Cross Gramian
- Empirical Cross Gramian
- Empirical Sensitivity Gramian
- Empirical Identifiability Gramian
- Empirical Joint Gramian



## Features:

- Interfaces for: Solver, inner product kernels & distributed memory
- Non-Symmetric option for all cross Gramians
- Compatible with OCTAVE and MATLAB
- Vectorized and parallelizable
- Open-source licensed

**More info:** <http://gramian.de>

## Single-Board Computers (SBC):

- Powerful CPU, but limited memory.
- Run Octave (on Linux).
- Low (per unit) cost.

## Model Reduction:

- Online phase (Evaluation) on SBC
- **Offline** phase (Assembly) on SBC
- Less model, less power (see PAC015<sup>12</sup>)

---

<sup>12</sup>C. Himpe and M. Ohlberger. **Zero-Copy Empirical Gramians**. Workshop on Power Aware Computing (PACO), 2015. [http://himpe.science/talks/himpe15\\_paco.pdf](http://himpe.science/talks/himpe15_paco.pdf)



## Hardware Test Subjects:

1. **Custom PC** (<http://products.amd.com/...>):

**CPU:** AMD APU (Kaveri / Steamroller)

**RAM:** 32GB DDR3

2. **HP EliteDesk 705 G3 Mini** (<http://www8.hp.com/...>):

**CPU:** AMD APU (Bristol Ridge / Excavator)

**RAM:** 32GB DDR4

3. **NanoPi NEO** (<http://nanopi.io/nanopi-neo.html>):

**CPU:** Allwinner ARM (Cortex v7)

**RAM:** 0.5GB DDR3

4. **NanoPi NEO2** (<http://nanopi.io/nanopi-neo2.html>):

**CPU:** Allwinner ARM (Cortex v8A)

**RAM:** 0.5GB DDR3



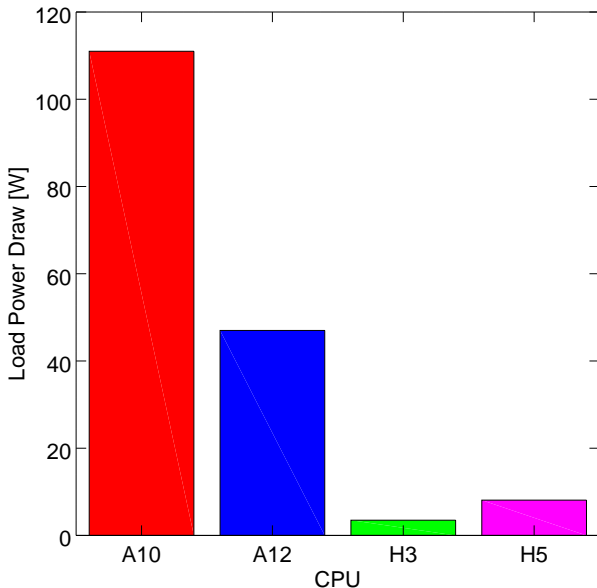
AMD	A10-7800	A12-9800E
Arch	x86-64(64-bit)	x86-64(64-bit)
Cores	4	4
Clock	3.5 Ghz	3.1 Ghz
Turbo	3.9 Ghz	3.8 Ghz
L1i-Cache	2 x 96 KB	2 x 96 KB
L1d-Cache	4 x 16 KB	<b>4 x 32 KB</b>
L2-Cache	<b>2 x 2 MB</b>	2 x 1 MB
SIMD	2x AVX+FMA3 4	2x <b>AVX2</b> +FMA3 4

- Performance vs size optimized routing
- Further instruction set extensions: MOVBE, BMI2, ...

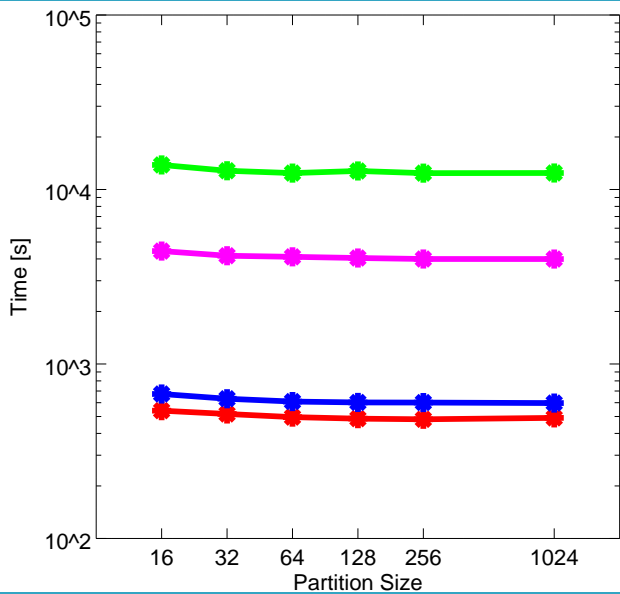


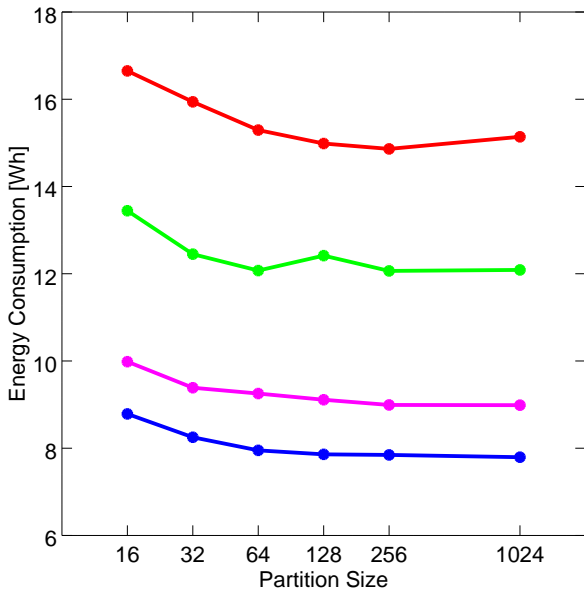
Allwinner	H3	H5
Arch	Cortex-A7(32-bit)	<b>Cortex-A53(64-bit)</b>
Cores	4	4
Clock	0.4 Ghz	0.4 Ghz
Turbo	0.8 Ghz	0.7 Ghz
L1i-Cache	4 x 32 KB	4 x 32 KB
L1d-Cache	4 x 32 KB	4 x 32 KB
L2-Cache	512 KB	512 KB
SIMD	NEON+VFP4	NEON+VFP4

- Instruction-Sets: V7 vs V8A
- Turbo set so temperature for same cooler is safe.









On the NanoPi NEO:

- Cooling is paramount!
- NEO2 can be pretty fast.
- Next: Mini cluster of 4+1 Pis.

Technicalities:

- Heuristic optimization optimization.
- `-O3` & `-ffast-math` (<http://youtu.be/w5Z4J1MJ1VQ>)
  - NEON SIMD requires `-ffast-math`.
- Abstract parallelization.
- A12 has an iGPU with 2 : 1 single-double ratio.



- The empirical cross Gramian is special.
- HAPOD allows low-rank SVD.
- SBCs are almost there.

Extendend Abstract: [doi:10.5281/zenodo.814498](https://doi.org/10.5281/zenodo.814498)

## Acknowledgment:

Supported by the German Federal Ministry for Economic Affairs and Energy, in the joint project: “**MathEnergy** – Mathematical Key Technologies for Evolving Energy Grids”, sub-project: Model Order Reduction (Grant number: 0324019**B**).