Exploring an Efficient Approach for Architecture-Level Thermal Simulation of Multi-core CPUs

Lin Jiang, Anthony Dowling, Yu Liu, Ming-C. Cheng
Department of Electrical & Computer Engineering
Clarkson University
Potsdam, NY, USA 13699
{jiangl2, dowlinah, yuliu, mcheng}@clarkson.edu

***Abstract*—In this work, an accurate and efficient thermal simulation approach based on proper orthogonal decomposition (POD) is applied to predict the temperature profile in an Intel Xeon E5-2699v3 CPU consisting of 18 cores. Using the POD method, the thermal problem is projected from a physical domain onto a functional space represented by a finite set of basis functions (or POD modes) that need to be trained by a large amount of thermal data. To generate the static and dynamic power consumption in space for the Xeon E5-2699v3 CPU as the heat source for thermal simulations in the training and validation, the cycle-level system simulator, gem5, and the power simulator, McPAT, are used. Gem5 is used to simulate the architectural characteristics of the CPU and gather performance counters. These statistics are gathered using a subset of the widely used SPLASH-2 benchmark suite as the simulated workload. These performance statistics contain information regarding architectural events such as functional unit usage, cache accesses. With the generated power trace, thermal data is collected using FEniCS, an open-source platform that supports finite element method (FEM)-based simulation, subjected to a range of power variations in space and time. It has been demonstrated that the proposed approach is able to offer an accurate thermal prediction of the CPU with a reduction in degrees of freedom (DoF) by more than 5 orders of magnitude and a speedup of 3 orders of magnitude, compared to FEM.**

***Index Terms*—Thermal simulation, proper orthogonal decomposition, data driven learning method, multi-core CPUs.**

I. Introduction

With the rapid miniaturization of integrated circuits (ICs) in the last several decades, high-performance microprocessors are becoming more thermally constrained due to the increasing power density [1], [2]. High temperature and hot spots resulting from joule heating seriously degrade the microprocessor performance and reliability [3]–[6]. To suppress hot spots and reduce processor temperature, thermal management and thermal-aware exploration can be implemented to decrease the possibility of CPU failure and improve processor performance. To achieve this effectively, accurate and efficient prediction of the thermal distribution in the processor is needed.

Different approaches have been developed for chip-level thermal simulation with different levels of accuracy and efficiency. These include the accurate but time-consuming direct numerical simulations (DNSs), such as the finite difference or finite element method (FDM or FEM, respectively), the more efficient Green's function method [7], [8], very efficient RC thermal circuits [9]–[12], data-driven learning algorithms [13], [14], etc. Each method has its own advantages and disadvantages in terms of complexity, assumptions, approximations, resolution, efficiency and accuracy.

In this work, we applied a data-driven learning method to dynamic thermal simulation of a multi-core CPU based on proper orthogonal decomposition (POD) [15], [16]. This POD thermal modeling technique offers an accurate and efficient prediction of the thermal profile in the processor without a priori assumptions. The POD projects the thermal problem in a physical domain of the processor onto a functional space, and its basis functions (or POD modes) for the space are extracted from thermal data generated by DNS of the CPU. In our study, the thermal data is collected from FEniCS, an open-source computing platform for solving partial differential equations (PDEs) using FEM [17]. To provide the heat sources more realistically, the power trace, including static and dynamic power, is generated by gem5 [18] and McPAT [19] with selected benchmarks. The validation of the POD approach was conducted using a different power trace than that used for thermal data collection. The results show that the POD approach is able to achieve a decrease in numerical degrees of freedom (DoF) by more than 5 orders of magnitude with a high accuracy, resulting in a speedup of 3 orders of magnitude, when compared with an FEM simulator.

II. Proper Orthogonal Decomposition for Thermal Simulation

POD is a projection-based reduced order model and able to represent a complex thermal problem in time and space with a small number of POD modes [20], [21]. These modes are extracted via a training process using the thermal solution data of the physical domain obtained by DNSs with a range of parametric variations including the boundary conditions (BCs) and dynamic power trace in our study. After the training

This work is supported by National Science Foundation under Grant No. ECCS-2003307.

978-1-6654-8485-5/22/$31.00 ©2022 IEEE

process, the POD modes are tailored to the range of the BCs and power variations.

A. Construction of POD Modes

The POD modes are optimized by maximizing the mean square inner product of the thermal solution data with the modes over the entire domain [15], [16], subjected to dynamic parametric variations of BCs and interior power sources. In our study, the thermal data in space is collected at each simulation time step in the DNS. This maximization process [15], [16] leads to an eigenvalue problem described by

$$\int_{\Omega} R(\vec{r},\vec{r}')\vec{\varphi}(\vec{r}')d\vec{r}' = \lambda\vec{\varphi}(\vec{r}), \tag{1}$$

where λ is the eigenvalue corresponding to the POD mode (i.e., eigenvector $\vec{\varphi}$) and $R(\vec{r},\vec{r}')$ is a two-point correlation tensor given as

$$R(\vec{r},\vec{r}') = \langle T(\vec{r},t) \otimes T(\vec{r}',t)\rangle \tag{2}$$

with $\otimes$ as the tensor operator and $\langle\rangle$ indicates the average over the number of thermal data sets. With the modes generated from the training process, including collecting the thermal data and solving the eigenvalue problem in (1), temperature can then be represented by

$$T(\vec{r},t) = \sum_{i=1}^{M} a_i(t)\varphi_i(\vec{r}), \tag{3}$$

where M is the number of POD modes selected to reconstruct the temperature solution. If the eigenvalue of the data decreases rapidly for the higher modes, only a small DoF (or M) is needed to reach an accurate thermal prediction.

B. Projection of the Thermal Problem to POD Space

To construct a POD model, the heat transfer equation is projected onto the POD space using the Galerkin projection

$$\int_{\Omega}(\varphi_i(\vec{r})\frac{\partial\rho CT}{\partial t} + \nabla\varphi_i \cdot k\nabla T)d\Omega = \int_{\Omega}\varphi_i(\vec{r})P_d(\vec{r},t)d\Omega - \int_{S}\varphi_i(\vec{r})(-k\nabla T\cdot\vec{n})dS, \tag{4}$$

where k, ρ and C are the thermal conductivity, density and specific heat, respectively. $P_d(\vec{r},t)$ is the interior power density, S is the boundary surface and $\vec{n}$ is the outward normal vector of the boundary surface. With the selected POD modes (4) can be rewritten as an M-dimensional ordinary differential equation (ODE) for $a_i(t)$,

$$\sum_{i=1}^{M} c_{i,j}\frac{da_i(t)}{dt} + \sum_{i=1}^{M} g_{i,j}a_i(t) = P_j,\ j = 1 \text{ to } M, \tag{5}$$

where $c_{i,j}$ and $g_{i,j}$ are the elements of thermal capacitance matrices and thermal conductance matrices in the POD space, respectively, and they are defined as

$$c_{i,j} = \int_{\Omega}\rho C\varphi_i\varphi_j d\Omega,\ g_{i,j} = \int_{\Omega} k\nabla\varphi_i\cdot\nabla\varphi_j d\Omega. \tag{6}$$

P_j in (5) is the power source strength for the j-th POD mode in the POD space and described as

$$P_j = \int_{\Omega}\varphi_j P_d(\vec{r},t)d\Omega - \int_{S}\varphi_j(\vec{r})(-k\nabla T\cdot\vec{n})dS. \tag{7}$$

Once the power consumption is obtained from gem5 and McPAT, the interior power source strength in POD space given in (7) can be pre-evaluated. For the boundary heat source in (7), the BC of the substrate bottom is modeled by convection heat transfer with a constant heat transfer coefficient and an ambient temperature of 45 °C (T_{amb}). All other boundaries are adiabatic. The coefficients in (6) can also be pre-evaluated once the modes are determined. With $a_i(t)$ solved from (5) in POD simulation, the temperature can then be predicted from (3).

III. Evaluation Methodology

The complete workflow of the efficient POD thermal simulation approach at the architecture level for multi-core CPUs is illustrated in Fig. 1. As shown in Fig. 1, the power trace is generated using two open-source simulators, gem5 and McPAT, before the FEniCS-FEM thermal simulation to collect thermal data. Gem5 was selected to perform cycle-level simulation of an Intel Xeon E5-2699v3 CPU and generate performance statistics of the CPU [18].

The SPLASH-2 benchmark suite was chosen as the simulated workload for gem5 [22]. SPLASH-2 is an open-source benchmark suites in C/C++, and it is widely used in the research field of computer architecture [23], [24]. 6 benchmarks are selected from this suite for POD training and validation.

The statistics traces output by gem5 for both the training and validation benchmarks are parsed to create XML input files for McPAT to simulate the power used by the CPU. McPAT also requires information regarding architectural information, such as the number of functional units, the size of caches. With the architectural information and performance statistics, McPAT calculates the dynamic and static power used by each component during each time step [19]. These power values are then parsed to create power trace files.

With the generated power trace, FEniCS-FEM is used to perform thermal simulation for data collection to generate POD modes as shown in Fig. 1. The trained POD modes then contain essential thermal information for the selected CPU subjected to a range of BCs and variations of the dynamic power trace via the training/data-collection process.

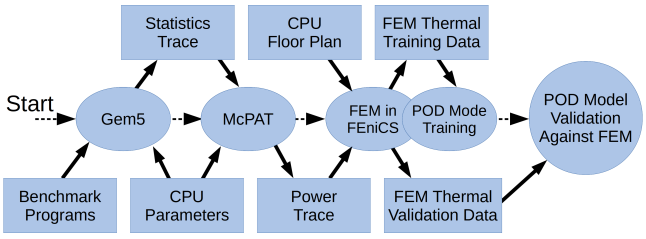

Fig. 1. Workflow of the POD thermal simulation approach for multi-core CPUs at the architecture level.

To validate the POD model, a separate thermal simulation for the selected CPU is performed using a different power trace. The POD results are then compared against the FEniCS-FEM simulation with consistent numerical settings and power trace.

IV. Implementing the POD Approach for Thermal Simulation of the Multi-core CPU

A. Data Collection from FEM Simulation

FEniCS-FEM is used to perform thermal simulation for the multi-core CPU, Intel Xeon E5-2699v3, whose floorplan is shown in Fig. 2 [25] to collect the temperature data for generation of the POD modes. The geometry of the multi-core CPU consists of a device layer and a substrate layer. The active areas and interconnects of the functional units are located in the device layer where power consumption occurs, and joule heating is generated. The generated heat in the CPU is mainly dissipated through the substrate via the heat spreader and heat sink. The BC of the substrate bottom is modeled by convection heat transfer with a constant heat transfer coefficient. This is modeled by a thermal resistance calculated from the chip dimension and material properties. All other surfaces are assumed adiabatic. The power density trace generated using gem5 and McPAT is uniform within each functional unit, and the dynamic power for each unit is averaged over 10k CPU cycles at 2.3 GHz. The data collection was carried in FEM simulations with meshes,$500 \times 347 \times 8$.

Using the collected temperature data, a set of POD modes and their eigenvalues were generated by solving the eigenvalue problem in (1) using the method of snapshots [20], [26]. These modes thus contain essential thermal information offered by the BCs and power trace used in the data collection in DNS. Each eigenvalue represents the mean squared temperature variations captured by each mode and thus reveals the significance of the mode.

The eigenvalue spectrum of the collected thermal data is illustrated in Fig. 3 for the first 25 modes. The eigenvalue decreases significantly in the first several modes and shows a slower decreasing rate in the high modes. The eigenvalue drops more than two orders of magnitude from the first to the second mode. A reduction by over three orders and nearly four orders of magnitude from the first to the third and fourth modes, respectively, is observed. It is therefore expected that POD

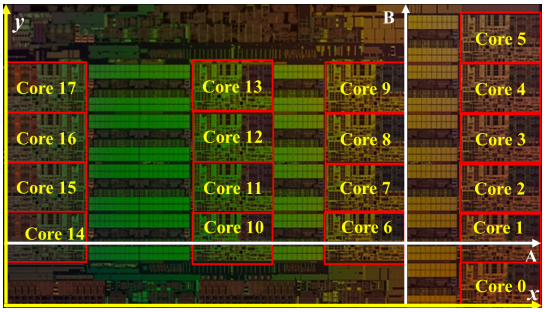

Fig. 2. Floorplan of the Intel Xeon E5-2699v3 CPU [25]. A and B indicate the plotting paths for the demonstration.

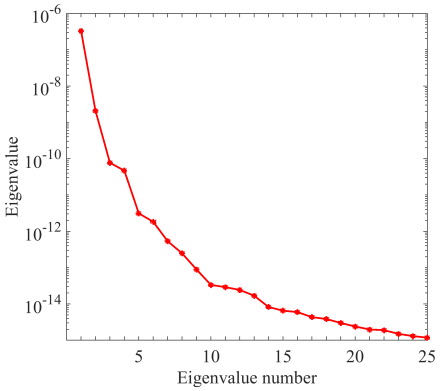

Fig. 3. Eigenvalue spectrum of the temperature data obtained from FEniCS-FEM simulation for the first 25 modes.

model for the multi-core CPU is able to offer a reasonably accurate thermal prediction with 3 or 4 modes.

B. Temperature Prediction for the Multi-core CPU from POD Thermal Simulation

To demonstrate the proposed POD approach, the model is used to perform thermal simulation of the selected multi-core CPU with identical BCs and a different power trace generated using a different assignment of benchmarks. The POD simulation result is compared against the FEniCS-FEM result to validate its accuracy because FEniCS-FEM was used to collect thermal data for the POD mode training. Computational efficiency of the developed POD model is also compared to the FEniCS-FEM in this work.

A comparison of the temperature evolution between POD model and FEM simulations at the location $(28\ mm, 4.4\ mm)$ is given in Fig. 4. The POD model with 3 or more POD modes is able to accurately predict temperature evolution. It is however interesting to observe a larger discrepancy between these 2 approaches with 3 modes for $0.2\ ms < t < 0.35\ ms$ and the deviation is effectively suppressed with more modes.

To illustrate the POD model accuracy in the entire CPU, temperature profiles at $2.6\ ms$ along Paths A and B are shown in Figs. 5 and 6, respectively. It can be seen that the temperature in the cores is higher than other components of the multi-core CPU due to the cores used by the running of benchmarks. In general, temporal and spatial thermal solutions derived from POD model, as presented in Figs. 4 - 6, indicate that the POD approach with 3 or more modes offers an accurate thermal prediction. Our study shows that a larger POD error is usually observed near locations in close proximity to very high thermal gradients due to less accurate numerical solution data used to train the POD modes. Use of smaller grids in the DNS simulation in general offers a better data quality and thus improves the effectiveness of the POD modes. This however significantly increases the computational cost for generating the POD modes and calculating the POD model parameters.

Overall, the deviation of the POD approach from the FEniCS-FEM reduces as more modes are included in the POD

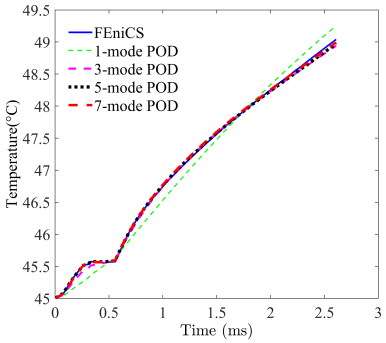

Fig. 4. Temperature evolution derived from POD model compared to FEniCS-FEM at $(28\ mm, 4.4\ mm)$.

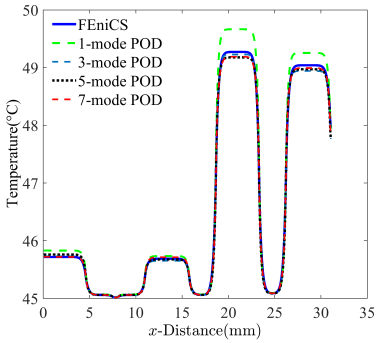

Fig. 5. The temperature distribution derived from POD model along Path A at 2.6 ms, compared to FEniCS-FEM.

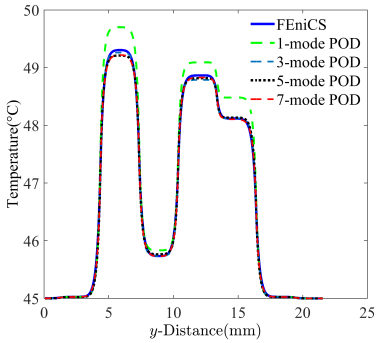

Fig. 6. The temperature distribution derived from POD model along Path B at 2.6 ms, compared to FEniCS-FEM.

simulation. This is achieved by the optimization process [15] that minimizes the least square (LS) error over the entire spatial and temporal domain. The LS error is given as

$$err_{ls} = \sqrt{\frac{\sum_{i=1}^{N_t} \int_\Omega e_i^2(\vec{r}) d\Omega}{\sum_{i=1}^{N_t} \int_\Omega (T_i(\vec{r}) - T_{amb})^2 d\Omega}}, \tag{8}$$

where $T_i(\vec{r})$ and $e_i(\vec{r})$ are the temperature solution given by FEM and the temperature difference between the FEM and the POD model at i-th time step, respectively, and N_t is the total number of time steps.

As shown in Table I, with 3 or more POD modes for the POD model an LS error near or below 3% can be reached, which offers a reduction in numerical DoF by more than 5 orders of magnitude ($500 \times 347 \times 8$ vs. 3), compared to the FEniCS-FEM simulation. Even using 7 modes with an LS error of 2.7%, the DoF reduction is still more than 5 orders of magnitude. When using more than 7 modes, the LS error will just reduce to a value slightly below 2.7% and become invariant with fluctuations.

TABLE I
LEAST SQUARE ERROR OF THE POD MODEL AGAINST FENICS-FEM.

Number of Modes	**1**	**3**	**5**	**7**
LS Error (%) of POD model	5.3	3.1	2.9	2.7

The demonstration of the POD modeling technique clearly illustrates the accuracy and efficiency for large-scale thermal simulation of the multi-core CPU. The efficiency enabled by a more-than 5 order reduction in DoF amounts to a significant saving in computational time. Comparison of computational times of FEniCS-FEM and POD model for simulation of the selected 18-core processor are shown in Table II. These times are collected on a Dell Precision Tower T7910, with dual Intel Xeon E5-2697A v4 CPUs, 512G memory, and Ubuntu 20.04 as its OS. It should be noted that the POD computational time listed in the Table II includes the time for solving $a_i(t)$ in (5) and the post processing time to recover temperatures in all 18 cores (instead of the entire chip) using (3). Actually, only temperature in some locations of interest is needed, which significantly reduces computing time and memory space needed in the post process. If the locations of higher power sources in the power trace are known, only temperature near those higher power sources are needed, which will further reduce the computing time and memory.

It can be seen in Tables I and II that the POD model with 3 modes, which offers an LS error near 3.1%, is 2056 times faster than the FEniCS-FEM. Even when 7 modes are used leading to an error near 2.7%, there is still a 1187 times of speedup. It is worth noting that the resolution of the POD solution is determined by its modes whose resolution is equivalent to the FEM used for their training.

TABLE II
COMPUTATIONAL TIMES OF FENICS-FEM AND POD MODEL.

Simulators	**FEniCS-FEM**	**Number of modes in POD model**		
		3	5	7
Time(s)	1686.7	0.82	1.14	1.42

V. CONCLUSION

A projection-based data-driven learning algorithm based on POD has been constructed for an 18-core CPU (Intel Xeon E5-2699v3) and used to investigate the dynamic thermal distribution in the processor. The FEniCS-FEM simulation was performed to collect thermal data for training the POD modes. The FEniCS-FEM was also used to validate the effectiveness of the proposed POD approach in terms of efficiency and accuracy. To apply more realistic static and dynamic power consumption in the selected processor, the SPLASH-2 benchmark suite was used to generate computing workload. The power trace including static and dynamic power, used as the heat source in FEM and POD simulations of the Xeon E5-2699v3 CPU was then obtained from the cycle-level system simulator, gem5, and the power simulator, McPAT, as described in Sec. III, based on the realistic workload.

It has been demonstrated that the developed POD model is able to provide an accurate prediction using a very small DoF. Compared to the FEniCS-FEM simulation, the POD approach offers a reduction in numerical DoF by more than 5 orders of magnitude. As a result, the POD model is able to achieve an LS error near 3% with 3 orders of magnitude speedup compared to the FEniCS-FEM simulation.

REFERENCES

[1] W. Jin, S. Sadiqbatcha, *et al.*, "Full-chip thermal map estimation for commercial multi-core cpus with generative adversarial learning," in *Proc. ICCAD*, pp. 1–9, 2020.

[2] S. I. Guggari, "Analysis of thermal performance metrics—application to cpu cooling in hpc servers," *IEEE Trans. Components, Packag. Manuf. Technol.*, vol. 11, no. 2, pp. 222–232, 2021.

[3] X. Zhou, Y. Xu, *et al.*, "Thermal management for 3d processors via task scheduling," in *Proc. ICPP*, pp. 115–122, 2008.

[4] H. F. Sheikh, I. Ahmad, Z. Wang, and S. Ranka, "An overview and classification of thermal-aware scheduling techniques for multi-core processing systems," *Sustain. Comput. Informat. Syst.*, vol. 2, no. 3, pp. 151–169, 2012.

[5] A. Heinig, R. Fischbach, and M. Dittrich, "Thermal analysis and optimization of 2.5 d and 3d integrated systems with wide i/o memory," in *Proc. ITHERM*, pp. 86–91, IEEE, 2014.

[6] J. Zhou, J. Yan, K. Cao, Y. Tan, T. Wei, M. Chen, G. Zhang, X. Chen, and S. Hu, "Thermal-aware correlated two-level scheduling of real-time tasks with reduced processor energy on heterogeneous mpsocs," *J. Syst. Archit.*, vol. 82, pp. 1–11, 2018.

[7] H. Sultan, A. Chauhan, and S. R. Sarangi, "A survey of chip-level thermal simulators," *ACM Comput. Surv.*, vol. 52, no. 2, pp. 1–35, 2019.

[8] S. Varshney, H. Sultan, P. Jain, *et al.*, "Nanotherm: An analytical fourier-boltzmann framework for full chip thermal simulations," in *Proc. ICCAD*, pp. 1–8, IEEE, 2019.

[9] "Hotspot 6.0 temperature modeling tool." http://lava.cs.virginia.edu/HotSpot. Accessed: 2021-05-21.

[10] A. Sridhar, A. Vincenzi, M. Ruggiero, *et al.*, "3d-ice: Fast compact transient thermal modeling for 3d ics with inter-tier liquid cooling," in *Proc. ICCAD*, pp. 463–470, IEEE, 2010.

[11] P. Chaparro, J. Gonzáles, G. Magklis, Q. Cai, and A. González, "Understanding the thermal implications of multi-core architectures," *IEEE Trans. Parallel Distrib. Syst.*, vol. 18, no. 8, pp. 1055–1065, 2007.

[12] E. Glocker and D. Schmitt-Landsiedel, "Modeling of temperature scenarios in a multicore processor system," *Adv. Radio Sci.*, vol. 11, no. D. 1, pp. 219–225, 2013.

[13] K. Zhang, A. Guliani, *et al.*, "Machine learning-based temperature prediction for runtime thermal management across system components," *IEEE Trans. Parallel Distrib. Syst.*, vol. 29, no. 2, pp. 405–419, 2017.

[14] A. Sridhar, A. Vincenzi, M. Ruggiero, and D. Atienza, "Neural network-based thermal simulation of integrated circuits on gpus," *IEEE Trans. Comput.-Aided Des.*, vol. 31, no. 1, pp. 23–36, 2012.

[15] J. L. Lumley, "The structure of inhomogeneous turbulent flows," *Atmospheric turbulence and radio wave propagation*, 1967.

[16] G. Berkooz, P. Holmes, and J. L. Lumley, "The proper orthogonal decomposition in the analysis of turbulent flows," *Annual review of fluid mechanics*, vol. 25, no. 1, pp. 539–575, 1993.

[17] "Fenics project." https://fenicsproject.org/. Accessed: 2021-05-21.

[18] N. Binkert, B. Beckmann, *et al.*, "The gem5 simulator," *SIGARCH Comp. Arch. News*, vol. 39, no. 2, pp. 1–7, 2011.

[19] S. Li, J. H. Ahn, *et al.*, "Mcpat: An integrated power, area, and timing modeling framework for multicore and manycore architectures," in *Proc. 42nd Ann. IEEE/ACM Int'l Symp. Microarchitecture*, pp. 469–480, 2009.

[20] W. Jia, B. T. Helenbrook, and M.-C. Cheng, "Fast thermal simulation of finfet circuits based on a multiblock reduced-order model," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 35, no. 7, pp. 1114–1124, 2016.

[21] W. Jia, B. T. Helenbrook, and M.-C. Cheng, "Thermal modeling of multi-fin field effect transistor structure using proper orthogonal decomposition," *IEEE Trans. Electron Devices*, vol. 61, pp. 2752–2759, 2014.

[22] S. C. Woo, M. Ohara, *et al.*, "The splash-2 programs: Characterization and methodological considerations," *SIGARCH Comp. Arch. News*, vol. 23, no. 2, pp. 24–36, 1995.

[23] J. Donald and M. Martonosi, "Power efficiency for variation-tolerant multicore processors," in *Proc. ISLPED*, pp. 304–309, 2006.

[24] G. Oxman and S. Weiss, "An noc simulator that supports deflection routing, gpu/cpu integration, and co-simulation," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 35, no. 10, pp. 1667–1680, 2016.

[25] H. Mujtaba, "Intel xeon e5-2600 v3 "haswell-ep" workstation and server processors unleashed for high-performance computing." https://wccftech.com/intel-xeon-e52600-v3-haswellep-workstation-server-processors-unleashed-highperformance-computing. Accessed: 2021-05-21.

[26] L. Sirovich, "Turbulence and the dynamics of coherent structures. i. coherent structures," *Quart. Appl. Math.*, vol. 45, pp. 561–571, 1987.