

AdEle: An Adaptive Congestion-and-Energy-Aware Elevator Selection for Partially Connected 3D NoCs

Ebadollah Taheri, Ryan G. Kim, and Mahdi Nikdast

Department of Electrical and Computer Engineering, Colorado State University, Fort Collins, CO 80523, USA

Abstract—By lowering the number of vertical connections in fully connected 3D networks-on-chip (NoCs), partially connected 3D NoCs (PC-3DNoCs) help alleviate reliability and fabrication issues. This paper proposes a novel, adaptive congestion- and energy-aware elevator-selection scheme called AdEle to improve the traffic distribution in PC-3DNoCs. AdEle employs an offline multi-objective simulated-annealing-based algorithm to find good elevator subsets and an online elevator selection policy to enhance elevator selection during routing. Compared to the state-of-the-art techniques under different real-application traffics and configuration scenarios, AdEle improves the network latency by 10.9% on average (up to 14.6%) with less than 6.9% energy consumption overhead.

Index Terms—Partially connected 3D networks-on-chip, through-silicon via, simulated annealing, elevator selection.

I. INTRODUCTION

Network-on-chip (NoC) has become the prevailing solution to enable scalable on-chip communication in manycore systems. Moreover, with the advances in three-dimensional (3D) integration technologies, 3D NoCs are emerging to further improve the heterogeneity and integration density by vertically stacking multiple dies connected with an efficient die-to-die interconnect [1]. Among different vertical interconnect technologies, through-silicon vias (TSVs) promise high bandwidth and low power [2]–[4].

TSV-based 3D NoCs have been proposed for various applications (e.g., [5], [6]). However, vertical links in TSV-based 3D NoCs use multiple TSVs in a bundle, resulting in high area overhead due to the large TSV interconnect-pitch and keep-out-zone requirements [7]. Also, TSVs are particularly susceptible to electromigration and capacitive crosstalk-induced issues [8], [9]. Therefore, 3D NoC architectures with TSVs at every router (i.e., fully connected) impose higher design complexity, fabrication costs, and performance degradation [1], [4]. Addressing such challenges has motivated the development of 3D NoCs with fewer TSV-based vertical links, also known as partially connected 3D NoCs (PC-3DNoCs) [1], [10], [11].

Nevertheless, PC-3DNoCs introduce some new design challenges because of their partial vertical connectivity [11], [12]. In particular, the vertical links (a.k.a. the elevators) must be shared among multiple routers, potentially creating traffic hotspots at the elevators and increasing the network latency [1]. To balance the traffic at these hotspot elevators, an adaptive routing technique is needed to select lower utilized elevators without detouring far from the minimal path (i.e., elevator-selection problem). Yet, initial routing solutions in PC-3DNoCs (e.g., Elevator-First routing [10]) naively select the nearest elevator without considering the traffic, resulting

in unbalanced elevator utilization. To reduce congestion, advanced methods (e.g., CDA [12]) use global traffic information to improve the traffic distribution during runtime. However, retrieving global traffic information increases both the hardware overhead and network traffic.

This paper addresses the elevator-selection problem in PC-3DNoC routing techniques by developing, for the first time, a novel, congestion- and energy-aware adaptive elevator-selection scheme called AdEle. AdEle works in two stages to balance the traffic with minimal overhead: an offline elevator-set optimization and an online elevator-selection policy. In the offline elevator-set optimization, AdEle uses a multi-objective simulated-annealing-based optimization algorithm (AMOSA [13]) to collectively choose an optimized subset of elevators for each source router that minimizes the average latency and energy under an assumed traffic scenario. During runtime, each router monitors its local traffic and selects one elevator from its subset to improve the latency of the network. AdEle employs a low-overhead local traffic monitoring technique that examines the blocking as a proxy for path congestion, balancing the elevator traffic while eliminating the overhead of global traffic monitoring used in other approaches. Our results simulated using different real-application traffics and configuration scenarios show the promise of AdEle compared to the state-of-the-art techniques: on average, AdEle improves the network latency by 10.9% (up to 14.6%) and with only 6% (up to 6.9%) energy consumption overhead.

The rest of the paper is organized as follows. We review the recent related work on PC-3DNoCs in Section II. Section III discusses the elevator-selection problem and its complexity in PC-3DNoCs and details our proposed technique and its implementation. We present our simulation results in Section IV. Finally, Section V concludes the paper.

II. BACKGROUND AND RELATED WORK

Employing conventional dimension-order routing algorithms in PC-3DNoCs will result in deadlock because of the irregular topology in such networks. To prevent deadlock, the Elevator-First routing algorithm [10] employs two virtual networks to break cyclic dependencies. Moreover, as the elevator-less routers cannot directly send packets to other layers, an elevator is selected for each packet to facilitate the inter-layer communication. Leveraging such a principle, several routing algorithms have been proposed for PC-3DNoCs [3], [14]. However, they follow an elevator-selection policy that ignores elevators' load distribution and the minimal path. This can be especially harmful for PC-3DNoCs with non-uniform elevator

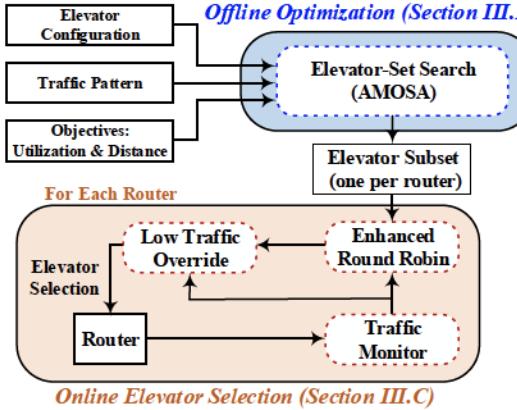


Fig. 1. An overview of our proposed elevator-selection scheme: AdEle.

placements, small number of elevators, or non-uniform traffic distributions. Adaptive elevator-selection techniques have been proposed [4], [11], [15] but mainly focus on elevator failure concerns. These strategies select the closest non-faulty elevator to the source without considering the elevator's congestion, causing them to suffer from high energy and latency costs.

To improve the traffic distribution in PC-3DNoCs, [16] proposed an optimized elevator-selection scheme using the Tabu search algorithm. However, the offline Tabu optimization cannot capture the dynamics of the runtime network traffic. Also, the search algorithm ignores the network energy efficiency during the elevator selection. In [12], an online elevator-selection scheme called CDA selects the elevator based on the buffer utilization of the routers between a source and the elevator. However, CDA requires online global information of the network buffer utilization which imposes high latency and hardware overheads to share this information.

Considering the aforementioned works, an efficient elevator-selection solution is essential but yet to be addressed for PC-3DNoCs. We take on this challenge by developing a novel, adaptive congestion- and energy-aware elevator-selection scheme (AdEle). Offline elevator-selection approaches enjoy low overhead while online approaches achieve better network latency and energy consumption. Accordingly, AdEle combines the benefits of both approaches while also considering energy consumption. On top of being energy-aware, AdEle includes elevator redundancy and online policies to accommodate dynamic traffic behavior. We will show that using a set of elevators instead of one elevator for each router can greatly improve network performance. Also, our proposed approach only utilizes local information of routers to effectively manage elevator congestion with low overheads.

III. PROPOSED ELEVATOR-SELECTION SCHEME: ADELE

This section details our proposed adaptive congestion- and energy-aware elevator-selection scheme. As shown in Fig. 1, AdEle uses an offline multi-objective simulated-annealing-based algorithm (AMOSA) to find an optimal subset of elevators for each router, and an online elevator-selection algorithm to improve elevator selection in the presence of runtime traffic. The following discusses the novel contributions of AdEle.

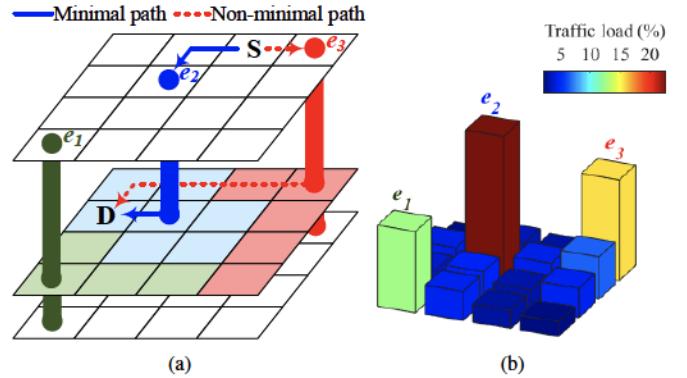


Fig. 2. (a) An example PC-3DNoC with three elevators (e_1 , e_2 , and e_3). The routing path from S to D based on Elevator-First algorithm [10] (dotted-red line) and the minimal path (blue-solid line) are shown. The middle-layer routers are colored based on their Elevator-First selected elevator. (b) Traffic load on each router in the middle layer: the e_2 elevator is highly congested because of the inefficient elevator selection in Elevator-First algorithm.

A. Motivation: Routing in PC-3DNoCs

In PC-3DNoCs, because of the irregular topology, the routing process requires three main steps: 1) selecting an elevator for each packet in the source router and then routing the packet to that elevator; 2) vertically routing the packet to the destination layer; and 3) routing the packet from the elevator to the destination node. In this routing process, the elevator selection (the first step) is critical as the number of vertical paths (elevators) is much smaller than the number of horizontal paths, putting significantly more traffic pressure on the elevators.

Fig. 2(a) shows an example of a PC-3DNoC with three elevators (e_1 – e_3) using Elevator-First-based elevator selection [4], [10], [11] (i.e., the closest elevator to the source router is selected). Routers are colored with the elevator's color they would use under the Elevator-First policy: i.e., four routers will use the green (e_1) elevator, seven will use the blue (e_2) elevator, and five will use the red (e_3) elevator. Unfortunately, such an uneven elevator utilization can put severe traffic pressure on certain elevators (e_2 in this example). Ideally, some of the load on the e_2 elevator could be assigned to the e_1 or e_3 elevators, making the e_2 elevator less congested. Fig. 2(b) demonstrates the utilization of the middle-layer routers with Elevator-First selection policy under uniform traffic. As can be seen, e_2 is highly congested due to the uneven elevator selection. In terms of energy efficiency, the best elevator selection is on the minimal path between the source and destination. However, as can be seen in Fig. 2(a) for the path between S and D, policies like Elevator-First (red-dotted line) do not necessarily choose the minimal path (blue-solid line).

AdEle will consider both traffic distribution and energy efficiency to select optimal elevators and evenly distribute traffic loads among the elevators. To the best of our knowledge, AdEle is the first congestion- and energy-aware elevator-selection scheme in PC-3DNoCs that includes elevator redundancy and online policies to accommodate dynamic traffic behavior while relying only on local router information.

B. Optimal Elevator-Subset for Each Router

To find the optimal subset of elevators for each router, AdEle performs an offline optimization to distribute the expected traffic load across all elevators and minimize the average inter-node (source to destination) distance. To do this, we first define two optimization objectives: 1) elevator-utilization variance to improve the traffic load distribution, and 2) average inter-node distance to minimize the energy consumption. Leveraging these objective functions, we will use a multi-objective simulated-annealing-based algorithm (AMOSA [13]) to find the optimal elevator subsets.

1) *Objective 1 - Elevator Utilization:* To balance the traffic on the elevators, AdEle attempts to minimize the elevator-utilization variance. As discussed above, it is important to evenly distribute the traffic over elevators to avoid highly congested elevators. To calculate the utilization variance, let us consider an N -node/router network with a set of elevators $\mathcal{E} = \{e_1, e_2, \dots, e_E\}$, where E is the total number of elevators. Moreover, assume that during runtime, each router i can select its elevator from a subset $A_i \subseteq \mathcal{E}$. For simplicity, for now we assume that each router selects each elevator from its elevator subset (A_i) uniformly (e.g., using a round-robin policy). Therefore, the utilization of elevator e (U_e) is:

$$U_e = \sum_{i=1}^N \frac{1}{|A_i|} \sum_{j=1}^N f_{ij} \cdot P_{ije}, \quad (1)$$

where f_{ij} is the frequency of traffic between routers i and j , and P_{ije} denotes whether the routing between routers i and j uses the elevator e ($P = 1$) or not ($P = 0$). Leveraging (1), the average traffic over all the elevators (μ) is:

$$\mu = \frac{1}{E} \sum_{i=1}^E U_i. \quad (2)$$

Using (1) and (2), elevator-utilization variance is:

$$\sigma^2 = \frac{1}{E} \sum_{i=1}^E (U_i - \mu)^2. \quad (3)$$

Minimizing the elevator-utilization variance will result in a better distribution of traffic load on the elevators and lower network latency.

2) *Objective 2 - Average Distance:* To improve network energy efficiency, AdEle attempts to minimize the average distance. As elevator selection is under consideration here, we only consider inter-layer traffic here. Therefore, the distance between inter-layer nodes i and j over an elevator e can be defined as:

$$D_{ij}^e = \begin{cases} 0, & i \text{ and } j \text{ are on the same layer} \\ d_{se} + d_e + d_{ed}, & \text{otherwise} \end{cases}, \quad (4)$$

where d_{se} , d_e , and d_{ed} are the Manhattan distances between the source and elevator, on the elevator (inter-die), and from the elevator to the destination, respectively. Based on (4), the average inter-layer-node distance in an L -layer network is:

$$AD = \frac{1}{N \times (\frac{L-1}{L} \times N)} \sum_{i=1}^N \frac{1}{|A_i|} \sum_{e=1}^{|A_i|} \sum_{j=1}^N D_{ij}^e. \quad (5)$$

3) *Multi-Objective Optimization:* We use a multi-objective simulated annealing-based optimization algorithm (AMOSA [13]) to find a set of optimal elevator subsets for all the routers in the network ($\mathcal{A} = \{A_1, \dots, A_N\}$) while minimizing the objective functions in (3) and (5). As AMOSA is a multi-objective optimization search, it offers a set of solutions that lie on the Pareto front of the optimization objectives (see [13] for more details). AMOSA-based optimization in AdEle provides different optimal solutions in terms of latency and energy efficiency. From these solutions, a designer can make trade-offs when choosing between more latency-aware or energy-aware solutions (see Fig. 3). Selection of solutions are discussed in detail in Section IV.

C. Adaptive Elevator Selection

Here, we discuss how a router i can efficiently select an elevator during runtime from its elevator subset (A_i) identified in the previous subsection. As we are interested in an even distribution of traffic load over all the elevators to improve traffic congestion during runtime, we apply an enhanced round-robin (RR) algorithm to select an elevator. In a conventional RR approach, elevators would be selected in a sequential order without considering the runtime traffic. To account for runtime traffic, we include the probability of skipping (P_{Sik}) a congested elevator (k) for router i in the RR approach. P_{Sik} is adjusted based on the average latency imposed by the elevator k , i.e., higher latencies seen using elevator k increases the probability of skipping it in the future. Accordingly, AdEle can adaptively manage dynamic traffic loads and congestion.

To find P_{Sik} , let us first define a cost function associated with making a selection from an elevator subset. After selecting an elevator, AdEle estimates the cost of this selection by considering the time between when the first flit (the header flit) and when the last flit (the tail flit) leave the source router. The latency (T_{ek}) imposed by selecting an elevator e_k from a subset A_i is:

$$T_{ek} = \frac{t_{tail} - t_{head} - l_p}{l_p}, \quad (6)$$

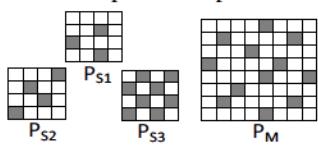
where t_{tail} and t_{head} denote the time when the tail flit and the header flit leave the source router, respectively. Also, l_p is the length of the packet. The elevator-selection cost (C_k) can be updated using the latency of the last selection defined in (6) and based on:

$$C_k \leftarrow (a \times T_{ek}) + ((1 - a) \times C_k), \quad 0 \leq a \leq 1 \quad (7)$$

where a is a coefficient to increase or decrease the impact of the new cost versus the old cost. We have experimentally found that $a = 0.2$ produces good results in AdEle.

Leveraging (7), AdEle can estimate the latency cost at the source router with *only local information*. With wormhole switching, any blocking in an elevator can be propagated along the path from the elevator to the source router. Therefore, blocking at a source router can be interpreted as blocking in the elevator. Note that incorporating global-network information into AdEle would improve the selection policy but be less practical as it will impose high hardware area, energy consumption, and latency costs.

TABLE I
SIMULATION SETUP

Network size	$4 \times 4 \times 4$ and $8 \times 8 \times 4$
Routing and VC selection (w/o elevator selection)	Elevator-First [10] (used to avoid deadlock)
Buffer depth	4 flits
Packet size	10–30 flits (random)
Traffic pattern	Uniform, Shuffle, and Real Elevator-placement patterns
	

Considering (7), we can define router i 's relative cost of selecting elevator k from A_i versus other possible elevators:

$$C_{ik}^{rel} = \frac{C_k}{\sum_{p=0}^{|A_i|} C_p}. \quad (8)$$

Based on the relative cost of a particular elevator selection, the possibility of skipping that elevator in the RR approach is:

$$P_{Sik} = \begin{cases} 1 - \xi, & \text{if } C_{ik}^{rel} \geq \frac{2}{N} \\ N \times (C_{ik}^{rel} - \frac{1}{N}) \times (1 - \xi), & \text{if } \frac{2}{N} > C_{ik}^{rel} \geq \frac{1}{N} \\ 0, & \text{otherwise} \end{cases} \quad (9)$$

Here, ξ is considered to allow for exploring new solutions even under high relative costs ($\xi = 0.05$ in our experiments). To clarify the use of ξ , suppose that the P_{Sik} of a selection is 1 because of high congestion. In this case, the elevator k will not be selected in the RR sequence at all and have no chance to update its elevator-selection cost (C_k). This would keep P_{Sik} high and prevent the elevator from observing any changes in its cost. To address such an update failure, ξ allows every elevator to be selected with a low probability regardless of P_{Sik} so the cost function has a chance for updating. To improve energy efficiency, when C_k is below a threshold for all k (low latency applications) and congestion is not a concern, AdEle will instead choose the elevator along the minimal path (discussed in Section III.A). Here, we experimentally find the threshold that minimizes the latency for each traffic and elevator configuration. Our future work will investigate a dynamic threshold management.

IV. SIMULATION AND EVALUATION RESULTS

Here, we compare AdEle against two well-known elevator-selection approaches: Elevator-First [10] and CDA [12]. The simulation setup is shown in Table I. In PC-3DNoCs, the number and location of elevators is limited by hardware constraints [9]. Therefore, AdEle is evaluated using different elevator-placement patterns to show that its efficacy is independent of any such patterns. Also, because of performance-area trade-off in PC-3DNoCs [1], various elevator concentrations might be employed. Therefore, here we simulate different concentration of elevators to show that AdEle performance is not limited by elevator concentration. Three elevator patterns are considered for a $4 \times 4 \times 4$ network (P_{S1} – P_{S3}) with different levels of elevator concentration. P_{S1} and P_{S3} are extracted to

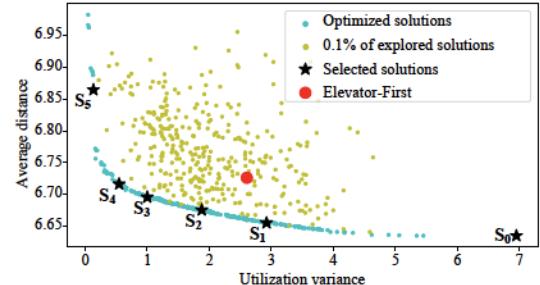


Fig. 3. Elevator-selection solutions found by AMOSA optimization in AdEle.

TABLE II
PERFORMANCE OF SELECTED SOLUTIONS FROM FIG. 3

	Elev. First	Optimized solutions				
		S_0	S_1	S_2	S_3	S_4
Latency*	161.4	396	209	156.6	76.9	67.4
Energy #	94.4	93.1	94.2	94.6	94.4	94.8

*Average Latency (cycles) # Energy/flit (n_j) \checkmark Selected

have an optimized average distance and P_{S2} is based on [4]. A large network ($8 \times 8 \times 4$) is simulated to show the scalability of AdEle. The pattern for this network (P_M) is also extracted based on the average distance optimization.

AdEle's offline optimization (see Section III.B) is implemented in Python to extract the elevator subsets for routers. These subsets are added to the AdEle router implemented in Access Noxim simulator [17]. We considered uniform traffic for the offline optimization, the most pessimistic assumption (i.e., traffic is not known *a priori*), while the network simulations are done using different synthetic and real-application traffics. Our analysis will demonstrate that AdEle does not require runtime traffic in its offline optimization as its online selection policy will adjust to runtime traffic. However, AdEle can use the runtime traffic during elevator-subset selection to offer further latency and energy improvement.

A. AMOSA Elevator-Subset Exploration

As discussed in Section III, AMOSA finds various solutions with different latency and energy-efficiency. To show the solution selection process, the optimization for P_M is detailed here. A small sample of AMOSA's explored solutions is shown in Fig. 3. As AMOSA explores the solution space, it makes its way towards the Pareto front (blue curve) to find the optimal trade-offs between utilization variance and average distance. Given the set of solutions, depending on the importance of energy efficiency (average distance) and latency (utilization variance), the final solution can be selected. For brevity, several of these points spread along the Pareto front are selected for network simulation (S_0 to S_5) where the results are summarized in Table II. Considering Table II and Fig. 3, lower utilization variance and lower average distance improves the latency and energy consumption, respectively. As we are able to significantly reduce the latency with fairly minimal increases in energy, we select S_5 for further analysis. Similarly, we select the solution for P_{S1} – P_{S3} .

B. AdEle Performance Under Synthetic Traffic

To compare AdEle with Elevator-First and CDA, we first evaluate the average latency under uniform and shuffle traf-

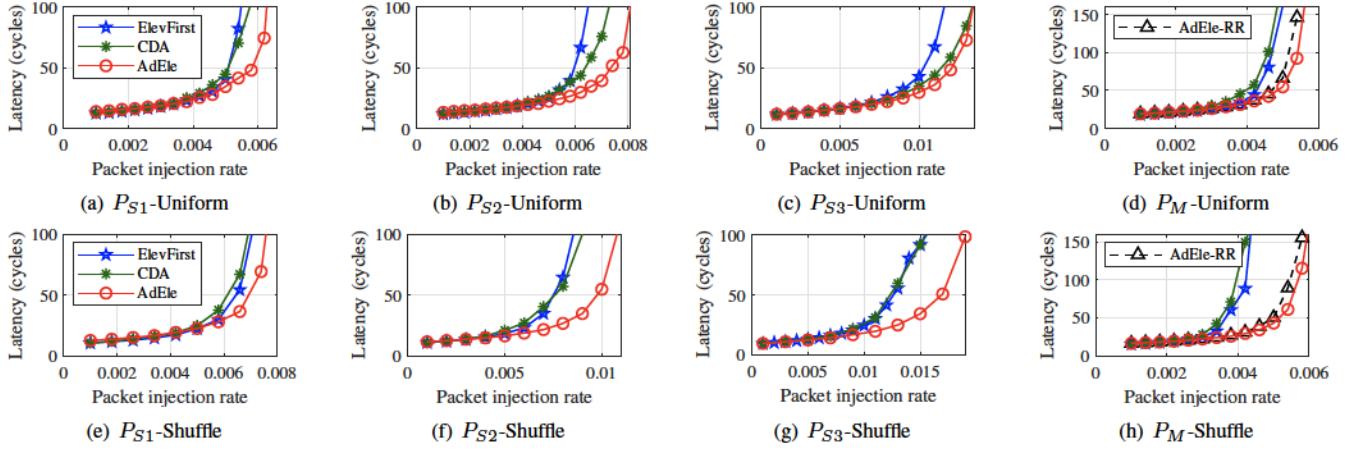


Fig. 4. Average latency for Elevator-First, CDA, and AdEle under uniform (a-d) and shuffle (e-h) traffic and with different elevator-placement patterns.

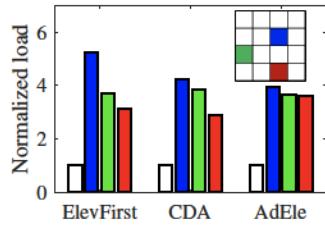


Fig. 5. Traffic load over routers with elevators (blue, green, and red) normalized to the average load over routers without an elevator (white bar).

fic patterns and with different elevator-placement patterns in Fig. 4. Across all the traffic and elevator-placement patterns, AdEle achieves the lowest latency and highest saturation threshold. Note that CDA is able to approach AdEle's performance because it considers global intra-layer traffic. In this work, we do not consider the high cost of CDA's global information sharing and optimistically assume that the information is instantaneously received at every router. In reality, CDA will likely perform much worse with stale information or include significant implementation overhead. With a higher elevator density (e.g., P_{S3}), the elevator congestion issue is less critical and intra-layer traffic will be more critical. Similarly, in a network with larger horizontal dimensions like P_M , intra-layer traffic is more important. Yet, AdEle shows better performance even with a high density of elevators and for P_M . Recall that AdEle's offline optimization step used uniform traffic. Yet, as Figs. 4(e)–(h) show, while the traffic is new for AdEle, it still achieves the lowest latency because its online selection policy can monitor runtime congestion and select better elevators. If the traffic is known *a priori*, AdEle can use this traffic information during offline optimization to improve elevator-subset selection even further. For P_M in Figs. 4(d) and 4(h), we also include the average latency of AdEle with standard RR selection. This demonstrates that AdEle's proposed online skipping policy achieves higher improvements in latency compared to RR in both uniform and shuffle traffic patterns.

To show the main reason for latency improvement when using AdEle, the load distribution over routers with elevators for P_{S1} is shown in Fig. 5. The white bar shows the average load over elevator-less routers. The other colored bars show the

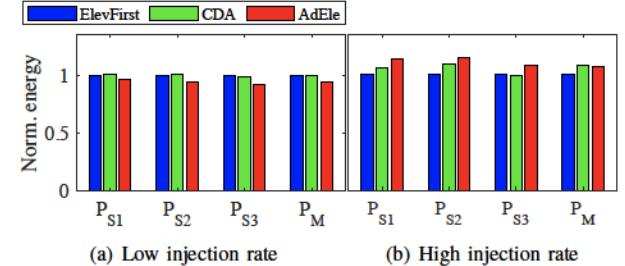


Fig. 6. Energy per flit for Elevator-First (ElevFirst), CDA, and AdEle normalized to ElevFirst, and under different injection rates.

load over different elevators. As can be seen, AdEle reduces the load on the highest utilized elevator (blue elevator). The energy consumption for each approach and elevator placement is shown in Fig. 6 for low ($1E-3$) and high ($5E-3$ to $1.2E-2$) injection rates based on the saturation point (injection rate at which latency is $10 \times$ zero-load latency) for each configuration. For low injection rates, AdEle has the lowest energy consumption because it switches to minimal routing and uses the minimal paths. On the other hand, AdEle incurs a small energy overhead (less than 9.7% compared to CDA) under high injection rates to take non-minimal paths and improve traffic congestion. If less energy overhead is desired, AdEle can use configurations with lower energy (see Table II).

C. AdEle Performance under Real-Application Traffic

We extracted the traffic of several SPLASH-2 [18] and PARSEC [19] benchmarks using Gem5 [20] for real-application simulations. Because Gem5 is limited to 64 cores, we demonstrate our results for P_{S1} – P_{S3} . As shown in Figs. 7(a)–(c), AdEle improves the network latency in nearly all cases. In particular, AdEle has more improvements in applications with higher traffic loads (canneal, fft, radix, and water) as there is more opportunity to reduce the resulting elevator congestion. In applications with lower traffic loads (fluidanimate and lu), AdEle maintains similar performance to the other approaches as there is little contention on the elevators and the latency is close to zero-load latency. Although P_{S1} still shows some improvements for AdEle, the lower number of elevators (three) results in minimal opportunity for AdEle to redirect traffic and

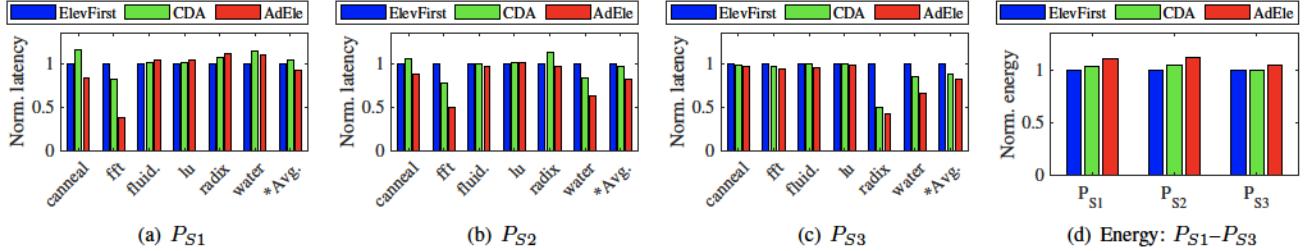


Fig. 7. Latency ((a)-(c), per application) and energy ((d)), averaged across all applications for Elevator-First (ElevFirst), CDA, and AdEle normalized to ElevFirst under real-application traffic with different elevator-placement patterns. *Avg. in (a)-(c) is the average of all six applications.

TABLE III
AREA ANALYSIS

	Cycles	Router area (μm^2)	Overhead
Base (ElevFirst)	1	35550	
CDA*	2	41088	14.4%
AdEle	1	36640	3.1%

*global information sharing is not included.

improve latency. On average, AdEle improves the network latency by 10.9% (up to 14.6%) compared to CDA and by 14.6% (up to 18%) compared to Elevator-First under P_{S1} – P_{S3} . Fig. 7(d) shows, for each elevator-placement pattern (P_{S1} – P_{S3}), the average energy over all the applications normalized to Elevator-First. AdEle imposes a small overhead because it may route packets over non-minimal paths in case of congestion to improve latency. Compared to CDA, AdEle has on average 6.9%, 6.2% and 4.8% energy overhead under P_{S1} , P_{S2} , and P_{S3} , respectively.

D. Hardware-Area Analysis and Comparison

Routers' hardware of Elevator-First, AdEle, and CDA are implemented and analyzed using Cadence Genus in 45 nm technology. Here, we consider a 1 GHz clock. The results are shown in Table III. Compared to CDA*, AdEle has a smaller area overhead. This is because AdEle only requires local traffic information while CDA requires a table to save global traffic information and find the best path in each router. However, CDA*'s area overhead is an optimistic assumption here as it does not include any overhead related to the actual sharing of information. Therefore, real CDA will likely impose higher area and latency overheads. Also, AdEle does not affect the router stages and will scale well with the network size, while CDA requires an additional cycle (or more for larger networks) to update its tables.

V. CONCLUSION

This paper proposes AdEle, an adaptive congestion- and energy-aware elevator-selection scheme to address elevator overutilization in partially connected 3D NoCs. Employing a set of elevators instead of one elevator for each source router, AdEle monitors the network traffic and provides an online policy to select the proper elevator while considering runtime traffic loads. AdEle only requires local router information and is able to improve average latency in various scenarios under both synthetic and real traffic at the cost of less than 6.9% in energy consumption. Moreover, AdEle can be easily adjusted to consider faults, which is of great interest in PC-3DNoCs, while considering elevator congestion.

ACKNOWLEDGMENT

This work was supported by the National Science Foundation (NSF) under grant number CNS-2046226.

REFERENCES

- [1] A. I. Arka *et al.*, "Making a case for partially connected 3D NoC: NFIC versus TSV," *ACM JETC*, vol. 16, no. 4, pp. 1–17, 2020.
- [2] T. Lu *et al.*, "TSV-based 3-D ICs: Design methods and tools," *IEEE TCAD*, vol. 36, no. 10, pp. 1593–1619, 2017.
- [3] R. Salamat *et al.*, "LEAD: An adaptive 3D-NoC routing algorithm with queuing-theory based analytical verification," *IEEE TC*, vol. 67, no. 8, pp. 1153–1166, 2018.
- [4] A. Coelho *et al.*, "FL-RuNS: A high-performance and runtime reconfigurable fault-tolerant routing scheme for partially connected three-dimensional networks on chip," *IEEE TNANO*, vol. 18, pp. 806–818, 2019.
- [5] X. Wang *et al.*, "HRC: A 3D NoC architecture with genuine support for runtime thermal-aware task management," *IEEE TC*, vol. 66, no. 10, pp. 1676–1688, 2017.
- [6] T. H. Vu *et al.*, "Fault-tolerant spike routing algorithm and architecture for three dimensional NoC-based neuromorphic systems," *IEEE Access*, vol. 7, pp. 90436–90452, 2019.
- [7] F. Wang *et al.*, "An effective approach of reducing the keep-out-zone induced by coaxial through-silicon-via," *IEEE T-ED*, vol. 61, no. 8, pp. 2928–2934, 2014.
- [8] T. Frank *et al.*, "Resistance increase due to electromigration induced depletion under TSV," in *IEEE IRPS*, 2011.
- [9] A. Eghbal *et al.*, "Analytical fault tolerance assessment and metrics for TSV-based 3D network-on-chip," *IEEE TC*, vol. 64, no. 12, pp. 3591–3604, 2015.
- [10] F. Dubois *et al.*, "Elevator-first: A deadlock-free distributed routing algorithm for vertically partially connected 3D NoCs," *IEEE TC*, vol. 62, no. 3, pp. 609–615, 2011.
- [11] E. Taheri *et al.*, "Addressing a new class of reliability threats in 3-D network-on-chips," *IEEE TCAD*, vol. 39, no. 7, pp. 1358–1371, 2020.
- [12] Y. Fu *et al.*, "Congestion-aware dynamic elevator assignment for partially connected 3D-NoCs," in *IEEE ISCAS*, 2019.
- [13] S. Bandyopadhyay *et al.*, "A simulated annealing-based multiobjective optimization algorithm: AMOSA," *IEEE Transactions on Evolutionary Computation*, vol. 12, no. 3, pp. 269–283, 2008.
- [14] J. Lee *et al.*, "Redef: An energy-efficient deadlock-free routing for 3D NoCs with partial vertical connections," *ACM JETC*, vol. 12, no. 3, pp. 1–22, 2015.
- [15] R. Salamat *et al.*, "A resilient routing algorithm with formal reliability analysis for partially connected 3D-NoCs," *IEEE TC*, vol. 65, no. 11, pp. 3265–3279, 2016.
- [16] S. Foroutan *et al.*, "Assignment of vertical-links to routers in vertically-partially-connected 3-D NoCs," *IEEE TCAD*, vol. 33, no. 8, pp. 1208–1218, 2014.
- [17] K.-Y. Jheng *et al.*, "Traffic-thermal mutual-coupling co-simulation platform for three-dimensional network-on-chip," in *IEEE VLSI-DAT*, 2010.
- [18] S. C. Woo *et al.*, "The SPLASH-2 programs: Characterization and methodological considerations," *ACM SIGARCH Computer Architecture News*, vol. 23, no. 2, pp. 24–36, 1995.
- [19] C. Bienia and K. Li, *Benchmarking modern multiprocessors*. Princeton University Princeton, NJ, 2011.
- [20] N. Binkert *et al.*, "The Gem5 simulator," *ACM SIGARCH Computer Architecture News*, vol. 39, no. 2, pp. 1–7, 2011.