

Communication-aware custom topology generation for VFI network-on-chip

Chang-Lin Li, Jae Hoon Lee, Joon-Sung Yang, and Tae Hee Han^{a)}

College of Information and Communication Engineering, Sungkyunkwan University, Suwon, Republic of Korea

a) than@skku.edu

Abstract: The voltage-frequency island (VFI) design paradigm has strong potential for achieving high energy efficiency in emerging communication centric many-core system-on-chip (SoC) technology called network-on-chip (NoC). In VFI design, each core is allocated into a suitable domain with different voltage and frequency considering the required workload. Most VFI NoC studies have focused on regular topologies. However, today's diversified applications demand custom NoC topologies to accommodate heterogeneous many-core architectures. In this paper, we propose a low-power custom topology generation method for constructing VFI NoCs, which embraces irregular topologies. The main objective of this work is to minimize energy consumption and communication resources, hence improving communication efficiency.

Keywords: voltage-frequency island, network-on-chip, custom topology generation, communication energy

Classification: Electron devices, circuits, and systems

References

- U. Y. Ogras, R. Marculescu, P. Choudhary and D. Marculescu: DAC 2007. Tech. Papers (2007) 110.
- [2] D. Mirzoyan, S. Stuijk, B. Akesson and K. Goossens: ESTIMedia 2013. Tech. Papers (2013) 1. DOI:10.1109/ESTIMedia.2013.6704497
- [3] J. Wang, Y. Qian, J. Lu and B. Li: ICISA 2014. Tech. Papers (2014) 1. DOI:10. 1109/ICISA.2014.6847386
- [4] B. Yu, S. Dong, S. Chen and S. Goto: ASP-DAC 2010. Tech. Papers (2010) 535. DOI:10.1109/ASPDAC.2010.5419825
- [5] K. Wang, S. Dong and S. Goto: GLSVLSI 2012. Tech. Papers (2012) 171. DOI:10.1145/2206781.2206823
- [6] D. Sengupta and R. Saleh: DAC 2008. Tech. Papers (2008) 155.
- [7] C. Seiculescu, S. Murali, L. Benini and G. D. Micheli: DAC 2009. Tech. Papers (2009) 822.
- [8] D. Shin, W. Kim, S. Kwon and T. H. Han: Des. Autom. Embed. Syst. 15 (2011) 89. DOI:10.1007/s10617-011-9070-x
- [9] A. B. Kahng, B. Li, L. S. Peh and K. Samadi: DATE 2009. Tech. Papers (2009) 423. DOI:10.1109/DATE.2009.5090700
- [10] N. Kapadia and S. Pasricha: GLSVLSI 2011. Tech. Papers (2011) 31. DOI:10. 1145/1973009.1973017





1 Introduction

The ever-increasing complexity of system-on-chips (SoCs), which is driven by Moore's law, is continually shifting the main design focus to on-chip interconnection networks to cope with communication bottlenecks. As the number of cores integrated on a single chip has grown remarkably, network-on-chip (NoC) technologies have been widely accepted as promising solutions for overcoming the bandwidth challenge of on-chip parallel communication architectures beyond shared bus and crossbar interconnections. Nevertheless, the faster growth rate of power dissipation relative to that of the power capacity or thermal design power (TDP) has become a critical challenge as the size of NoC scales up.

In order to mitigate this problem, each core should be assigned to an optimal voltage and frequency, which will provide a better power-performance tradeoff than the single voltage/frequency design counterpart. However, assigning different voltage/frequency values to all cores requires excessive mixed clock FIFOs (mcFIFOs), voltage level converters (VLCs), power rails, and voltage regulators, which adversely affects both power consumption and chip area. The voltage-frequency island (VFI) technique is a compromising solution for achieving system-level power management, where several cores are assigned to a common voltage/frequency [1].

A number of studies have proposed low-power VFI NoC designs based on mapping and partitioning techniques [2, 3]. Despite the significant potential for energy savings when using VFIs, most VFI NoC studies have focused on regular topologies. In particular, 2D mesh topology-based NoC architecture is commonly used. However, recently developed leading-edge SoCs consist of many heterogeneous cores with different computing and communication properties. Therefore, regular topologies are not sufficient to accommodate recent architectural trends. On the contrary, custom topologies can provide more design optimization opportunities by exploiting the asymmetric properties in both heterogeneous many-core SoC architectures and their associated communication traffic patterns [4, 5].

In [6], the authors present an application-driven and floorplan-aware design process for voltage islands, which categorizes cores according to their power consumption in order to allocate cores into appropriate voltage domains. However, they do not address communication power reduction.

A synthesis-based approach for customized NoC, which supports the power shutdown of voltage islands was suggested in [7], but router optimization which is crucial in VFI NoC topology generation was not assessed.

In this paper, we propose a low-power custom topology generation method for VFI NoC considering communication volume, core voltage domain, and NoC components such as routers and mcFIFOs. Both inter- and intra-VFI communication are analyzed to minimize the overall energy consumption by exploiting the characteristics of the custom topology. The contributions of this work are twofold. First, an energy model for the VFI NoCs using a custom topology by incorporating independent intra- and inter-communication is proposed. Second, the topology generation algorithm minimizes communication energy through three phases, i)





VFI-aware core clustering, ii) *network component generation*, and iii) *VFI-aware routing path allocation*. To the best of our knowledge, this is the first work on custom topology generation for VFI NoC.

2 Problem formulation

NoC is a highly structured and scalable solution, which comprises cores and network components. Data transactions between cores are managed through a router that consists of a crossbar switch, switch allocator, input/output buffers, and routing computation module. In VFI NoC, communication across different VFIs requires mcFIFOs and VLCs. The topology must ensure that each router can be reached from all of the other routers on the network. However, the number of mcFIFOs and VLCs in VFI NoC affects performance/energy and design complexities.

In regular topology-based VFI NoCs, each router is connected to a core, as shown in Fig. 1(a), where the communication between two cores requires more than two routers. On the other hand, the cores within an island may be connected to the same router or communicate over the other routers in custom topology-based VFI NoC, as shown in Fig. 1(b), thereby enabling us to reduce the number of routers and mcFIFOs. Furthermore, custom topology can provide expanded design space to meet design goals in terms of area, performance, and power consumption for emerging SoC architectures. To maximize the advantages of using custom topologies, the communication traffic characteristics of applications should be incorporated.

In this work, communication-aware custom topology generation starts from core clustering to minimize inter-VFI communication, which is closely related to the number of mcFIFOs and network components. Next, the routing paths are allocated to realize the minimum number of inter-VFI links. In order to formulate the problem, we introduce some definitions.

Definition 1: A core communication graph, CCG(V, E), where vertex $v_i \in V$ represents a core with an assigned voltage, and each directed edge $e_{i,j} \in E$ represents the communication volume from v_i to v_j .

Definition 2: A VFI NoC topology, N(C, R, S), where $c_i \in C$ represents the set of cores, $r_i \in R$ represents the set of routers, and $s_i \in S$ represents the set of voltage/ frequency islands. The number of VFIs, M, determines the cost of the mcFIFO.

Using the above definitions, the problem of custom topology generation for VFI NoCs can be formulated as follows.

Given a CCG(V, E), the number of islands, and a set of routers with p ports, generate N(C, R, S) such that allocate cores into an appropriate island and routing path to minimize the communication energy, E_{total} , in the VFI NoC.

The total communication energy in a VFI NoC (E_{total}) can be decomposed into inter-VFI communication energy (E_{inter}) and intra-VFI communication energy (E_{intra}) as follows:









Fig. 1. VFI NoC based on (a) regular topology, and (b) custom topology.

$$E_{total} = E_{inter} + E_{intra} \tag{1}$$

$$E_{inter} = \sum_{k=1}^{1} \left(E_{mcFIFO} \times n_k^{inter} + E_L \times \sum_{i=1}^{M} n_{i,k}^{intra} + E_R \times \sum_{i=1}^{M} (n_{i,k}^{intra} - 1) \right)$$
(2)

$$E_{intra} = \sum_{k=1}^{Z} \left(E_L \times \sum_{i=1}^{M} n_{i,k}^{intra} + E_R \times \sum_{i=1}^{M} (n_{i,k}^{intra} - 1) \right)$$
(3)

where E_{mcFIFO} is the energy consumption in the mcFIFO, n_k^{inter} is the number of inter-VFI links that the flits have to pass through from source to destination with respect to the *k*-th inter-VFI communication, E_L and E_R are the energy consumption in the link and router within each VFI, respectively, $n_{i,k}^{intra}$ is the number of intra-VFI links in the *i*-th VFI with respect to the *k*-th intra-VFI communication, *T* is the number inter-VFI communication, and *Z* is the number of intra-VFI communication.

3 Communication-aware VFI NoC topology generation algorithm

In order to minimize the energy required for communication, the proposed algorithm generates a topology considering the communication volume, core voltage domain, and NoC component. Because the cost of the inter-VFI is larger than that of the intra-VFI, we focus on reducing inter-VFI communications when





Input: $CCG(V, E)$ and the number of VFIs; Result: $N(C, R, S)$	
Result : $N(C, R, S)$ Phase 1: VFI-aware core clustering 1: sort the cores according to the $voltage(V)$ and find the lowest $voltage(V_L)$ and the highest $voltage(V_H)$ 2: calculate $\delta V = (V_H - V_L) / M$ 3: allocate cores to VFIs such that satisfy $(V < k \cdot \delta V, k = 1,, M)$ 4: set the voltage/frequency value in each VFI to the maximum one among the ones within the same VFI 5: calculate current energy consumption 6: rank inter-VFI communication in descending order 7: for each inter-VFI communication do 8: migrate a core to the VFI of the other core and calculate temporary energy 9: if temporary energy < current energy 10: update cluster 11: else 12: try core-pair with the next highest communication 13: end Phase 2: Network component generation 14: place a router for each VFI initially 15: for each intra-VFI communication do 16: find the pair of cores with larger communication 17: if remaining ports exist 18: connect the cores to current router 19: else 20: generate	Input : $CCG(V, E)$ and the number of VFIs;
Phase 1: VFI-aware core clustering1: sort the cores according to the $voltage(V)$ and find the lowest $voltage(V_L)$ and the highest $voltage(V_H)$ 2: calculate $\delta V = (V_H - V_L) / M$ 3: allocate cores to VFIs such that satisfy $(V < k \cdot \delta V, k = 1,, M)$ 4: set the voltage/frequency value in each VFI to the maximum one among the ones within the same VFI5: calculate current energy consumption6: rank inter-VFI communication in descending order7: for each inter-VFI communication do8: migrate a core to the VFI of the other core and calculate temporary energy9: if temporary energy < current energy	Result : $N(C, R, S)$
Phase 1: VFI-aware core clustering1: sort the cores according to the $voltage(V)$ and find the lowest $voltage(V_L)$ and the highest $voltage(V_H)$ 2: calculate $\delta V = (V_H - V_L) / M$ 3: allocate cores to VFIs such that satisfy $(V < k \cdot \delta V, k = 1,, M)$ 4: set the voltage/frequency value in each VFI to the maximum one among the ones within the same VFI5: calculate current energy consumption6: rank inter-VFI communication in descending order7: for each inter-VFI communication do8: migrate a core to the VFI of the other core and calculate temporary energy9: if temporary energy < current energy10: update cluster11: else12: try core-pair with the next highest communication 13: end14: place a router for each VFI initially15: for each intra-VFI communication do16: find the pair of cores with larger communication 17: if remaining ports exist18: connect the cores to current router19: else20: generate a new router and connect to cores21: end22: generate NI between connected routers and cores21: for each inter-VFI communication do23: for each inter-VFI links on the minimal routing path 25: for each inter-VFI link do26: associated link.cnt++27: end28: end29: select the inter-VFI link with the largest link.cnt	
1: <i>sort</i> the cores according to the <i>voltage(V)</i> and find the <i>lowest voltage(V_L)</i> and the <i>highest voltage(V_H)</i> 2: <i>calculate</i> $\delta V = (V_H - V_L) / M$ 3: <i>allocate</i> cores to VFIs such that satisfy ($V < k \cdot \delta V$, $k = 1,, M$) 4: <i>set</i> the voltage/frequency value in each VFI to the maximum one among the ones within the same VFI 5: <i>calculate current energy</i> consumption 6: <i>rank</i> inter-VFI communication in descending order 7: for <i>each</i> inter-VFI communication do 8: <i>migrate</i> a core to the VFI of the other core and calculate <i>temporary energy</i> < <i>current energy</i> 9: if <i>temporary energy < current energy</i> 10: <i>update</i> cluster 11: else 12: <i>try</i> core-pair with the next highest communication 13: end Phase 2: <i>Network component generation</i> 14: <i>place</i> a router for each VFI initially 15: for <i>each</i> intra-VFI communication do 16: <i>find</i> the pair of cores with larger communication 17: if remaining <i>ports exist</i> 18: <i>connect</i> the cores to current router 19: else 20: <i>generate</i> a new router and connect to cores 21: end 22: <i>generate</i> NI between connected routers and cores Phase 3: <i>VFI-aware routing path allocation</i> 23: for <i>each</i> inter-VFI links on the minimal routing path 25: for <i>each</i> inter-VFI link do 26: associated <i>link.cnt</i> ++ 27: end 28: end 29: <i>select</i> the inter-VFI link with the largest <i>link.cnt</i>	Phase 1: VFI-aware core clustering
<i>lowest voltage</i> (V_L) and the <i>highest voltage</i> (V_{th}) 2: <i>calculate</i> $\delta V = (V_H - V_L) / M$ 3: <i>allocate</i> cores to VFIs such that satisfy ($V < k \cdot \delta V$, $k = 1,, M$) 4: <i>set</i> the voltage/frequency value in each VFI to the maximum one among the ones within the same VFI 5: <i>calculate current energy</i> consumption 6: <i>rank</i> inter-VFI communication in descending order 7: for <i>each</i> inter-VFI communication do 8: <i>migrate</i> a core to the VFI of the other core and calculate <i>temporary energy</i> 9: if <i>temporary energy</i> < <i>current energy</i> 10: <i>update</i> cluster 11: else 12: <i>try</i> core-pair with the next highest communication 13: end Phase 2: <i>Network component generation</i> 14: <i>place</i> a router for each VFI initially 15: for <i>each</i> intra-VFI communication do 16: <i>find</i> the pair of cores with larger communication 17: if remaining <i>ports exist</i> 18: <i>connect</i> the cores to current router 19: else 20: <i>generate</i> a new router and connect to cores 21: end 22: <i>generate</i> NI between connected routers and cores Phase 3: <i>VFI-aware routing path allocation</i> 23: for <i>each</i> inter-VFI communication do 24: <i>find</i> all inter-VFI links on the minimal routing path 25: for <i>each</i> inter-VFI link do 26: associated <i>link.cnt</i> ++ 27: end 28: <i>end</i> 29: <i>select</i> the inter-VFI link with the largest <i>link.cnt</i>	1: sort the cores according to the $voltage(V)$ and find the
 2: calculate δV = (V_H - V_L) / M 3: allocate cores to VFIs such that satisfy (V < k · δV, k = 1,, M) 4: set the voltage/frequency value in each VFI to the maximum one among the ones within the same VFI 5: calculate current energy consumption 6: rank inter-VFI communication in descending order 7: for each inter-VFI communication do 8: migrate a core to the VFI of the other core and calculate temporary energy 9: if temporary energy < current energy 10: update cluster 11: else 12: try core-pair with the next highest communication 13: end Phase 2: Network component generation 14: place a router for each VFI initially 15: for each intra-VFI communication do 16: find the pair of cores with larger communication 17: if remaining ports exist 18: connect the cores to current router 19: else 20: generate a new router and connect to cores 21: end 22: generate NI between connected routers and cores Phase 3: VFI-aware routing path allocation 23: for each inter-VFI communication do 24: find all inter-VFI link son the minimal routing path 25: for each inter-VFI link do 26: associated link.cnt++ 27: end 28: end 29: select the inter-VFI link with the largest link.cnt 	<i>lowest</i> voltage(V_L) and the highest voltage(V_H)
 3: allocate cores to VFIs such that satisfy (V < k · δV, k = 1,, M) 4: set the voltage/frequency value in each VFI to the maximum one among the ones within the same VFI 5: calculate current energy consumption 6: rank inter-VFI communication in descending order 7: for each inter-VFI communication do 8: migrate a core to the VFI of the other core and calculate temporary energy 9: if temporary energy < current energy 10: update cluster 11: else 12: try core-pair with the next highest communication 13: end Phase 2: Network component generation 14: place a router for each VFI initially 15: for each intra-VFI communication do 16: find the pair of cores with larger communication 17: if remaining ports exist 18: connect the cores to current router 19: else 20: generate a new router and connect to cores 21: end 22: generate NI between connected routers and cores Phase 3: VFI-aware routing path allocation 23: for each inter-VFI communication do 24: find all inter-VFI links on the minimal routing path 25: for each inter-VFI link do 26: associated link.cnt++ 27: end 28: end 29: select the inter-VFI link with the largest link.cnt	2: calculate $\delta V = (V_H - V_L) / M$
 , M) set the voltage/frequency value in each VFI to the maximum one among the ones within the same VFI calculate current energy consumption rank inter-VFI communication in descending order for each inter-VFI communication do migrate a core to the VFI of the other core and calculate temporary energy if temporary energy < current energy if temporary energy < current energy update cluster else try core-pair with the next highest communication for each intra-VFI communication do find the pair of cores with larger communication if remaining ports exist connect the cores to current router generate a new router and connect to cores generate a new router and connect to cores for each inter-VFI communication do for each inter-VFI communication do for each inter-VFI communication do for each the cores to current router else generate a new router and connect to cores for each inter-VFI communication do for each inter-VFI communication do generate NI between connected routers and cores 	3: allocate cores to VFIs such that satisfy $(V < k \cdot \delta V, k =$
 4: set the voltage/frequency value in each VFI to the maximum one among the ones within the same VFI 5: calculate current energy consumption 6: rank inter-VFI communication in descending order 7: for each inter-VFI communication do 8: migrate a core to the VFI of the other core and calculate temporary energy 9: if temporary energy < current energy 10: update cluster 11: else 12: try core-pair with the next highest communication 13: end Phase 2: Network component generation 14: place a router for each VFI initially 15: for each intra-VFI communication do 16: find the pair of cores with larger communication 17: if remaining ports exist 18: connect the cores to current router 19: else 20: generate a new router and connect to cores 21: end 22: generate NI between connected routers and cores Phase 3: VFI-aware routing path allocation 23: for each inter-VFI links on the minimal routing path 25: for each inter-VFI link do 26: associated link.cnt++ 27: end 28: end 29: select the inter-VFI link with the largest link.cnt	1,, <i>M</i>)
one among the ones within the same VFI 5: calculate current energy consumption 6: rank inter-VFI communication in descending order 7: for each inter-VFI communication do 8: migrate a core to the VFI of the other core and calculate temporary energy 9: if temporary energy < current energy 10: update cluster 11: else 12: try core-pair with the next highest communication 13: end Phase 2: Network component generation 14: place a router for each VFI initially 15: for each intra-VFI communication do 16: find the pair of cores with larger communication 17: if remaining ports exist 18: connect the cores to current router 19: else 20: generate a new router and connect to cores 21: end 22: generate NI between connected routers and cores Phase 3: VFI-aware routing path allocation 23: for each inter-VFI communication do 24: find all inter-VFI links on the minimal routing path 25: for each inter-VFI link do 26: associated link.cnt++ 27: end 28: end 29: select the inter-VFI link with the largest link.cnt	4: set the voltage/frequency value in each VFI to the maximum
 5: calculate current energy consumption 6: rank inter-VFI communication in descending order 7: for each inter-VFI communication do 8: migrate a core to the VFI of the other core and calculate temporary energy 9: if temporary energy < current energy 10: update cluster 11: else 12: try core-pair with the next highest communication 13: end Phase 2: Network component generation 14: place a router for each VFI initially 15: for each intra-VFI communication do 16: find the pair of cores with larger communication 17: if remaining ports exist 18: connect the cores to current router 19: else 20: generate a new router and connect to cores 21: end 22: generate NI between connected routers and cores Phase 3: VFI-aware routing path allocation 23: for each inter-VFI link do 26: associated link.cnt++ 27: end 28: end 29: select the inter-VFI link with the largest link.cnt	one among the ones within the same VFI
 6: rank inter-VFI communication in descending order 7: for each inter-VFI communication do 8: migrate a core to the VFI of the other core and calculate temporary energy 9: if temporary energy < current energy 10: update cluster 11: else 12: try core-pair with the next highest communication 13: end Phase 2: Network component generation 14: place a router for each VFI initially 15: for each intra-VFI communication do 16: find the pair of cores with larger communication 17: if remaining ports exist 18: connect the cores to current router 19: else 20: generate a new router and connect to cores 21: end 22: generate NI between connected routers and cores Phase 3: VFI-aware routing path allocation 23: for each inter-VFI links on the minimal routing path 25: for each inter-VFI link do 26: associated link.cnt++ 27: end 28: end 29: select the inter-VFI link with the largest link.cnt	5: <i>calculate current energy</i> consumption
 7: for each inter-VFI communication do 8: migrate a core to the VFI of the other core and calculate temporary energy 9: if temporary energy < current energy 10: update cluster 11: else 12: try core-pair with the next highest communication 13: end Phase 2: Network component generation 14: place a router for each VFI initially 15: for each intra-VFI communication do 16: find the pair of cores with larger communication 17: if remaining ports exist 18: connect the cores to current router 19: else 20: generate a new router and connect to cores 21: end 22: generate NI between connected routers and cores Phase 3: VFI-aware routing path allocation 23: for each inter-VFI communication do 24: find all inter-VFI links on the minimal routing path 25: for each inter-VFI link do 26: associated link.cnt++ 27: end 28: end 29: select the inter-VFI link with the largest link.cnt	6: <i>rank</i> inter-VFI communication in descending order
 <i>migrate</i> a core to the VFI of the other core and calculate <i>temporary energy</i> <i>if temporary energy < current energy</i> <i>update</i> cluster <i>else</i> <i>try</i> core-pair with the next highest communication <i>end</i> <i>Phase 2: Network component generation Phase a router for each VFI initially for each</i> intra-VFI communication do <i>else cornect the cores to current router else generate</i> a new router and connect to cores <i>else generate</i> NI between connected routers and cores <i>Phase 3: VFI-aware routing path allocation 23: for each</i> inter-VFI communication do <i>24: find</i> all inter-VFI inks on the minimal routing path <i>25: for each</i> inter-VFI link do <i>26: associated link.cnt++ 27: end 28: end 29: select</i> the inter-VFI link with the largest <i>link.cnt</i>	/: IOP each inter-VFI communication do
9: if temporary energy < current energy	8: <i>migrate</i> a core to the VFI of the other core and calculate
 9: If temporary energy < current energy 10: update cluster 11: else 12: try core-pair with the next highest communication 13: end Phase 2: Network component generation 14: place a router for each VFI initially 15: for each intra-VFI communication do 16: find the pair of cores with larger communication 17: if remaining ports exist 18: connect the cores to current router 19: else 20: generate a new router and connect to cores 21: end 22: generate NI between connected routers and cores Phase 3: VFI-aware routing path allocation 23: for each inter-VFI communication do 24: find all inter-VFI links on the minimal routing path 25: for each inter-VFI link do 26: associated link.cnt++ 27: end 28: end 29: select the inter-VFI link with the largest link.cnt 	temporary energy
 10. <i>update</i> cluster 11: else 12: <i>try</i> core-pair with the next highest communication 13: end Phase 2: Network component generation 14: place a router for each VFI initially 15: for each intra-VFI communication do 16: <i>find</i> the pair of cores with larger communication 17: if remaining ports exist 18: <i>connect</i> the cores to current router 19: else 20: generate a new router and connect to cores 21: end 22: generate NI between connected routers and cores Phase 3: VFI-aware routing path allocation 23: for each inter-VFI communication do 24: <i>find</i> all inter-VFI links on the minimal routing path 25: for each inter-VFI link do 26: associated <i>link.cnt++</i> 27: end 28: end 29: select the inter-VFI link with the largest <i>link.cnt</i> 	9: In temporary energy < current energy
 11. erse 12: try core-pair with the next highest communication 13: end Phase 2: Network component generation 14: place a router for each VFI initially 15: for each intra-VFI communication do 16: find the pair of cores with larger communication 17: if remaining ports exist 18: connect the cores to current router 19: else 20: generate a new router and connect to cores 21: end 22: generate NI between connected routers and cores Phase 3: VFI-aware routing path allocation 23: for each inter-VFI communication do 24: find all inter-VFI links on the minimal routing path 25: for each inter-VFI link do 26: associated link.cnt++ 27: end 28: end 29: select the inter-VFI link with the largest link.cnt 	10. <i>update</i> cluster
 <i>Phase 2: Network component generation</i> 14: <i>place</i> a router for each VFI initially 15: for <i>each</i> intra-VFI communication do 16: <i>find</i> the pair of cores with larger communication 17: if remaining <i>ports exist</i> 18: <i>connect</i> the cores to current router 19: else 20: <i>generate</i> a new router and connect to cores 21: end 22: <i>generate</i> NI between connected routers and cores Phase 3: VFI-aware routing path allocation 23: for <i>each</i> inter-VFI communication do 24: <i>find</i> all inter-VFI links on the minimal routing path 25: for <i>each</i> inter-VFI link do 26: associated <i>link.cnt++</i> 27: end 28: end 29: <i>select</i> the inter-VFI link with the largest <i>link.cnt</i> 	11. CISC
Phase 2: Network component generation 14: place a router for each VFI initially 15: for each intra-VFI communication do 16: find the pair of cores with larger communication 17: if remaining ports exist 18: connect the cores to current router 19: else 20: generate a new router and connect to cores 21: end 22: generate NI between connected routers and cores Phase 3: VFI-aware routing path allocation 23: for each inter-VFI communication do 24: find all inter-VFI links on the minimal routing path 25: for each inter-VFI link do 26: associated link.cnt++ 27: end 28: end 29: select the inter-VFI link with the largest link.cnt	13: end
Phase 2: Network component generation 14: place a router for each VFI initially 15: for each intra-VFI communication do 16: find the pair of cores with larger communication 17: if remaining ports exist 18: connect the cores to current router 19: else 20: generate a new router and connect to cores 21: end 22: generate NI between connected routers and cores 23: for each inter-VFI communication do 24: find all inter-VFI links on the minimal routing path 25: for each inter-VFI link do 26: associated link.cnt++ 27: end 28: end 29: select the inter-VFI link with the largest link.cnt	13. Chu
Phase 2: Network component generation 14: place a router for each VFI initially 15: for each intra-VFI communication do 16: find the pair of cores with larger communication 17: if remaining ports exist 18: connect the cores to current router 19: else 20: generate a new router and connect to cores 21: end 22: generate NI between connected routers and cores Phase 3: VFI-aware routing path allocation 23: for each inter-VFI communication do 24: find all inter-VFI links on the minimal routing path 25: of each inter-VFI link do 26: associated link.cnt++ 27: end 28: end 29: select the inter-VFI link with the largest link.cnt	
 14: place a router for each VFI initially 15: for each intra-VFI communication do 16: find the pair of cores with larger communication 17: if remaining ports exist 18: connect the cores to current router 19: else 20: generate a new router and connect to cores 21: end 22: generate NI between connected routers and cores Phase 3: VFI-aware routing path allocation 23: for each inter-VFI communication do 24: find all inter-VFI links on the minimal routing path 25: for each inter-VFI link do 26: associated link.cnt++ 27: end 28: end 29: select the inter-VFI link with the largest link.cnt	Phase 2: Network component generation
 15: for each intra-VFI communication do 16: find the pair of cores with larger communication 17: if remaining ports exist 18: connect the cores to current router 19: else 20: generate a new router and connect to cores 21: end 22: generate NI between connected routers and cores Phase 3: VFI-aware routing path allocation 23: for each inter-VFI communication do 24: find all inter-VFI links on the minimal routing path 25: for each inter-VFI link do 26: associated link.cnt++ 27: end 28: end 29: select the inter-VFI link with the largest link.cnt	14: <i>place</i> a router for each VFI initially
 <i>find</i> the pair of cores with larger communication <i>if</i> remaining <i>ports exist</i> <i>connect</i> the cores to current router <i>else</i> <i>generate</i> a new router and connect to cores <i>end</i> <i>generate</i> NI between connected routers and cores <i>Phase 3: VFI-aware routing path allocation Si for each</i> inter-VFI communication do <i>for each</i> inter-VFI links on the minimal routing path <i>for each</i> inter-VFI link do <i>select</i> the inter-VFI link with the largest <i>link.cnt</i>	15: for each intra-VFI communication do
 17: if remaining <i>ports exist</i> 18: <i>connect</i> the cores to current router 19: else 20: <i>generate</i> a new router and connect to cores 21: end 22: <i>generate</i> NI between connected routers and cores Phase 3: VFI-aware routing path allocation 23: for each inter-VFI communication do 24: <i>find</i> all inter-VFI links on the minimal routing path 25: for each inter-VFI link do 26: associated <i>link.cnt</i> ++ 27: end 28: end 29: <i>select</i> the inter-VFI link with the largest <i>link.cnt</i>	16: <i>find</i> the pair of cores with larger communication
 18: connect the cores to current router 19: else 20: generate a new router and connect to cores 21: end 22: generate NI between connected routers and cores 23: for each inter-VFI communication do 24: find all inter-VFI links on the minimal routing path 25: for each inter-VFI link do 26: associated link.cnt++ 27: end 28: end 29: select the inter-VFI link with the largest link.cnt 	17: if remaining <i>ports exist</i>
 19: else 20: generate a new router and connect to cores 21: end 22: generate NI between connected routers and cores Phase 3: VFI-aware routing path allocation 23: for each inter-VFI communication do 24: find all inter-VFI links on the minimal routing path 25: for each inter-VFI link do 26: associated link.cnt++ 27: end 28: end 29: select the inter-VFI link with the largest link.cnt 	18: <i>connect</i> the cores to current router
 20: generate a new router and connect to cores 21: end 22: generate NI between connected routers and cores Phase 3: VFI-aware routing path allocation 23: for each inter-VFI communication do 24: find all inter-VFI links on the minimal routing path 25: for each inter-VFI link do 26: associated link.cnt++ 27: end 28: end 29: select the inter-VFI link with the largest link.cnt 	19: else
 21: end 22: generate NI between connected routers and cores <i>Phase 3: VFI-aware routing path allocation</i> 23: for each inter-VFI communication do 24: find all inter-VFI links on the minimal routing path 25: for each inter-VFI link do 26: associated link.cnt++ 27: end 28: end 29: select the inter-VFI link with the largest link.cnt 	20: <i>generate</i> a new router and connect to cores
 22. generate INI between connected rotters and cores Phase 3: VFI-aware routing path allocation 23: for each inter-VFI communication do 24: find all inter-VFI links on the minimal routing path 25: for each inter-VFI link do 26: associated link.cnt++ 27: end 28: end 29: select the inter-VFI link with the largest link.cnt 	21: end 22: consume NI between connected residers and comes
Phase 3: VFI-aware routing path allocation 23: for each inter-VFI communication do 24: find all inter-VFI links on the minimal routing path 25: for each inter-VFI link do 26: associated link.cnt++ 27: end 28: end 29: select the inter-VFI link with the largest link.cnt	22: generate INI between connected routers and cores
Phase 3: VFI-aware routing path allocation23: for each inter-VFI communication do24: find all inter-VFI links on the minimal routing path25: for each inter-VFI link do26: associated link.cnt++27: end28: end29: select the inter-VFI link with the largest link.cnt	
 23: for each inter-VFI communication do 24: find all inter-VFI links on the minimal routing path 25: for each inter-VFI link do 26: associated link.cnt++ 27: end 28: end 29: select the inter-VFI link with the largest link.cnt 	Phase 3: VFI-aware routing path allocation
 24: <i>find</i> all inter-VFI links on the minimal routing path 25: for <i>each</i> inter-VFI link do 26: associated <i>link.cnt++</i> 27: end 28: end 29: <i>select</i> the inter-VFI link with the largest <i>link.cnt</i> 	23: for each inter-VFI communication do
 25: for each inter-VFI link do 26: associated link.cnt++ 27: end 28: end 29: select the inter-VFI link with the largest link.cnt 	24: find all inter-VFI links on the minimal routing path
 26: associated <i>link.cnt</i>++ 27: end 28: end 29: <i>select</i> the inter-VFI link with the largest <i>link.cnt</i> 	25: for <i>each</i> inter-VFI link do
27: end28: end29: <i>select</i> the inter-VFI link with the largest <i>link.cnt</i>	26: associated <i>link.cnt</i> ++
28: end29: <i>select</i> the inter-VFI link with the largest <i>link.cnt</i>	27: end
29: <i>select</i> the inter-VFI link with the largest <i>link.cnt</i>	28: end
	29: <i>select</i> the inter-VFI link with the largest <i>link.cnt</i>
30: generate mcFIFO and VLC for each inter-VFI link;	30: generate mcFIFO and VLC for each inter-VFI link;

Fig. 2. Pseudo-code of the proposed algorithm.

allocating cores to VFIs. Given the input information, clusters having the same voltage and frequency are generated first. Next, routers and network interfaces (NIs) are created for each VFI. Finally, the VFI-aware routing path allocation is performed to build the minimum number of inter-VFI links to minimize the total communication cost.

Because the topology generation is known as an NP-hard problem, we adopted a heuristic algorithm that consists of the following three phases: i) *VFI-aware core clustering*, ii) *network component generation*, and iii) *VFI-aware routing path allocation*. The pseudo-code of our topology generation algorithm is described in Fig. 2.





In the first phase, all cores are sorted in order of decreasing supply voltage level, and δV is then calculated by dividing the maximum voltage difference by the number of VFIs (line 1–2). Each core is allocated to a VFI according to its voltage (line 3). The voltage/frequency value of all cores in the same VFI is set to the maximum one among the cores (line 4). The communication energy is estimated by assuming one- and two-hop distances for inter- and intra-VFI communication, respectively (line 5). This assumption is needed to generate an initial topology, and the final value of the hop distance will be determined in later steps. Iteratively, the pair of cores with the larger inter-VFI communication is selected, and each core is checked to determine if the total energy consumption can be reduced by moving it to the other VFI. Otherwise, the core pair with the next highest communication cost is considered if none of the alternative core migrations can reduce the total energy consumption (line 7–13). Accordingly, the inter-VFI links are transformed into intra-VFI links for the migrated core-pairs. The algorithm stops when there is no improvement in energy consumption after subsequent rearrangement of VFIs.

In the second phase, routers are placed to build up an initial VFI NoC topology (line 14). Considering the increased design complexity associated with the number of ports in routers, we employed a typical 4-port router, which is usual in 2D NoC design. Next, we performed intra-VFI optimization to place routers in a VFI. For each intra-VFI communication, the core-pair with larger communication volume is connected to the same router if there is an available port in the router (line 15–21). Finally, NIs are generated between the routers and cores (line 22).

In the third phase, we allocated routing paths to minimize the number of inter-VFI links, thus reducing the total communication cost. For each inter-VFI communication, the inter-VFI links on the minimal routing path is searched and each associated *link.cnt* is incremented (lines 23–28). Among the candidate inter-VFI links in each VFI communication, the link with the largest *link.cnt* indicates the most frequently used one. Thus, the inter-VFI link with the largest *link.cnt* is selected to minimize the number of inter-VFI links (line 29). Finally, mcFIFOs and VLCs are generated for the chosen optimal inter-VFI links (line 30).

4 Experimental results

The proposed method was evaluated using four traffic-intensive multimedia benchmarks, namely VOPD, MWD, MPEG4, and D_38, as shown in Table I. Using the four benchmarks, the proposed method (*prop*) was compared to the mesh topologybased VFI NoC (*topo1*) [8], floorplan-aware method (*topo2*) [6], and shutdown of the voltage island-based method (*topo3*) [7].

Table I. Specification of bench	marks.
---------------------------------	--------

	Number of cores	Number of edges	Average commu- nication volume [flits]
VOPD	12	14	227
MWD	12	12	470
MPEG4	12	13	521
D_38	38	46	247





Orion 2.0 [9] was used to estimate the energy consumption of the router and link with different voltages and frequencies using the parameter of 65 nm technology. The supply voltage level was varied from 0.7 V to 1.4 V in 0.1 V unit, and the operating frequency was varied from 200 MHz to 500 MHz in 50 MHz unit. We assume that the energy consumption in mcFIFO and VLC is 20% of that of the connected router, as in [10].



Fig. 3. Energy consumption comparison.

Fig. 3 shows the experimental results for the relative energy consumption with respect to *topo*1. With the exception of VOPD, the proposed method exhibits the best result. The average communication volume of VOPD is the smallest of the four benchmarks with the same number of cores, and the proposed method in VOPD is therefore not fully utilized with regard to the minimized inter-VFI links. The improvement in the average energy consumption relative to *topo*1, *topo*2, and *topo*3 are 23.02%, 9.29%, and 4.92%, respectively. In the case of *topo*3, the number of generated routers is less than that of the proposed method. However, *topo*3 produces routers with variable ports, which results in an increased energy consumption in routers. Therefore, the energy consumption in *topo*3 will increase more rapidly than that of the proposed one as the number of cores will increase.

In addition, speedup is estimated with our SystemC-based cycle-accurate simulator to evaluate the performance effect of the proposed method as shown in Table II. For all benchmarks, the proposed method shows the best result due to the use of customized topologies and application-specific communication traffic patterns.

Table II. Relative speedup results when using the proposed technique.

	vs. <i>topo</i> 1	vs. topo2	vs. topo3
VOPD	1.15	1.07	1.15
MWD	1.2	1.13	1.06
MPEG4	1.21	1.07	1.14
D_38	1.12	1.05	1.06





5 Conclusion

In this paper, we proposed a communication-aware custom topology generation method for VFI NoC to achieve higher energy efficiency. We incorporated extended topology beyond regular ones and application-specific traffic patterns, which are greatly needed in state-of-the-art communication-centric heterogeneous many-core SoCs. The three-phase heuristic algorithm exploits the characteristics of applications and custom topologies to minimize the communication energy. Simulation results demonstrated the effectiveness of the proposed method in terms of energy and performance.

Acknowledgments

This work was supported by MSIP (Ministry of Science, ICT & Future Planning), Korea, under ITRC (Information Research Center) support program supervised by the NIPA (National IT Industry Promotion Agency) (NIPA-2014-H0301-14-1018) and Samsung Research Fund.

