

Hierarchical Power Monitoring for On-chip Networks

Liang Guang, Alexander Wei Yin, Pekka Rantala, Ethiopia Nigusie, Pasi Liljeberg,
Jouni Isoaho, Hannu Tenhunen

Department of Information Technology, University of Turku, Finland
{liagua, yinwei, peaura, ethnig, pakrli, jisoaho, hatenhu}@utu.fi

1. Introduction

On-chip networks are constantly expanding in their sizes with transistor scaling, and thousand-core multi-processor is expected in the foreseeable future [1]. While sharing some similarities with their off-chip counterparts, on-chip networks (or NoCs) are critically limited by the power constraints [2]. A hierarchical power monitoring design flow and system architecture is presented in this paper. By implementing a hierarchy of adaptive agents, various power-monitoring services are handled with different granularities in a scalable and adaptive manner. Simulation results so far demonstrate the superior efficiency to existing architectures. With all functions realized in the future, flexible ultra-low-power management can be expected for large-sized NoCs.

2. Hierarchical power monitoring Design Flow

Hierarchical agent monitoring platform (initially presented in our previous work [3]), which provides a general architecture for a wide range of monitoring services, can be utilized for ultra-low-power monitoring. The agent hierarchy is composed of application agent, platform agent, cluster agent and cell agent. The application agent dynamically captures the application requirement including processing speed and power constraints at run-time. The platform agent provides system level monitoring, including resource mapping and network configuration to achieve minimal energy and fault tolerance. The cluster agent monitors the resources in each cluster (comprised of a number of cores with accessory switches, links and possibly caches or memories) with finer-granular services. Each cell agent provides the most fine-grained control of each processor, switch or link. Application and platform agents are being developed as software which can provide diverse functions, while cluster and cell agents will include dedicated hardware circuits which

provide single but fast local control. Fig. 1 illustrates how the agent hierarchy can be mapped on a regular mesh-based on-chip network (PE stands for processing element; NI stands for network interface; Sw stands for switch).

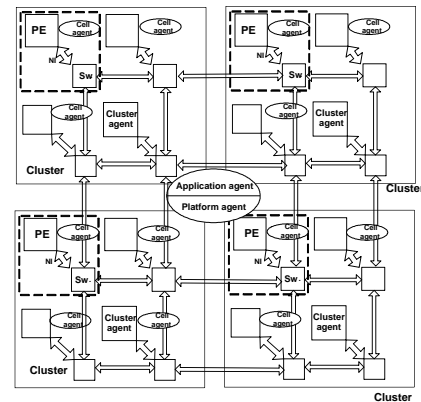


Fig. 1. Hierarchical agents on NoC structure

Based on the agent hierarchy, a low-power design flow is developed (Fig. 2). It starts with clusterization of the network into a number of clusters. The clusterization algorithm takes the inputs including the estimation of agent implementation overhead and communication overhead, and determines an optimal number of clusters (assuming regular cluster division) as well as the location of cluster and platform agents (cell agents are always located close to the components being monitored).

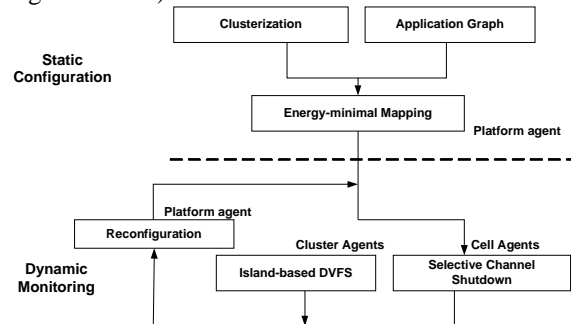


Fig. 2. Adaptive Hierarchical Power Monitoring Design Flow

The application, represented by its communication computation graph, will then be mapped onto processors with the total communication energy minimized. These static configuration steps are all performed by the platform agent. Despite that the overhead of calculating energy-minimal configuration can be large, it is incurred exclusively to the single platform agent on any-sized network and used only before execution or when the system requires reconfiguration.

At run-time, cluster-level DVFS (dynamic voltage frequency scaling) and cell-level power gating are performed adaptively. For cluster-level DVFS, each cluster is configured as a supply island, equipped with a DC converter and a PLL. The average network load in the cluster will be used to adjust the cluster voltage and frequency. The cluster-level DVFS is a scalable tradeoff between single domain and per-core supply scaling mechanisms. Only one DC converter and one PLL are required for a whole cluster. Each cell agent controls a dedicated power switch circuit (with sleep transistors), which temporarily shuts down the local channel when the network load is small. This fine-grained power gating can greatly reduce leakage consumption with minimal overhead (the control and switching of sleep transistors) and controllable performance penalty. In case of severe functional failures or the application requires resource remapping, system reconfiguration will be handled by the platform agent. Fig. 3 illustrates a segment of the NoC platform with hierarchical low-power monitoring (FIFO which supports multiple frequency operations, is needed to interface two supply domains. Cell agents are implicitly located adjacent to each core. The platform agent is omitted.)




Fig. 3. NoC equipped with Hierarchical Power Monitoring

3. Simulation result up-to-date

Cluster-based DVFS has been analyzed on a transport-level NoC simulator targeting 65nm technology, using a number of synthetic traffic patterns categorized by the injection and distribution variation patterns (injection patterns can be linearly changing or

b-model. Source pattern can be uniform or with hotspots. Destination pattern can be locality or with hotspots. Details are omitted in this paper). The network has 15*15 nodes regularly divided into 25 clusters, with three discrete supply levels (1.3V, 1.2G; 0.75V, 0.8G; 0.6V, 0.6G). Table 1 compares the energy consumption of the cluster-based DVFS with two other low-energy architectures: single-domain DVFS and static voltage-island configuration. The switch energy is obtained from Orion simulator, and the link energy is obtained from Cadence. The voltage scaling overhead is calculated based on [4].

Table 1. Normalized Energy Consumption for Three Low-Energy Architectures

Traffic pattern	Cluster-based DVFS	Centralized DVFS	Static voltage-island
1	80.90%	106.29%	1
2	78.12%	95.48%	1
3	79.36%	101.98%	1
4	96.21%	100.41%	1
5	73.20%	86.13%	1
6	90.18%	106.52%	1

From Table 1, we can observe that cluster-based monitoring can better exploit the temporal and spatial variation in minimizing the communication energy. As we are currently simulating, when combined with adaptive channel shutdown (power-gating), further leakage power minimization can be achieved.

4. Conclusion & ongoing work

The hierarchical power monitoring approach aims to develop a design flow and system architecture to provide ultra-low-power/energy for future NoCs in a configurable and scalable manner. We are currently simulating the agent algorithms at each level with estimation of implementation overhead.

References

- [1] Shekhar Borkar. Thousand core chips: a technology perspective. In *Proc. 44th ACM/IEEE DAC '07*, pages 746–749. 2007.
- [2] J. D. Owens, W. J. Dally, R. Ho, D. N. Jayasimha, S. W. Keckler, and Li-Shiuan Peh. Research challenges for on-chip interconnection networks. *IEEE MICRO*, 27(5):96–108, Sept.–Oct. 2007.
- [3] Pekka Rantala, Jouni Isoaho, and Hannu Tenhunen. Novel agent-based management for fault-tolerance in network-on-chip. In *Proc. of DSD 2007*, pages 551–555, 2007.
- [4] Anthony John Stratakos. *High-efficiency low-voltage DC-DC conversion for portable applications*. PhD thesis, University of California, Berkeley, 1998.