INTEGRATION, the VLSI journal 42 (2009) 479–485

Contents lists available at ScienceDirect

INTEGRATION, the VLSI journal

journal homepage: www.elsevier.com/locate/vlsi

A variable frequency link for a power-aware network-on-chip (NoC)

Seung Eun Lee *, Nader Bagherzadeh

Department of EECS, University of California-Irvine, Irvine, CA 92697, USA

ARTICLE INFO

Article history:
Received 2 February 2008
Received in revised form
22 January 2009
Accepted 25 January 2009

PACS:
85.40.e

Keywords:
Interconnection network
Network-on-chip (NoC)
Dynamic frequency scaling
Power optimization
System-on-chip (SoC)

ABSTRACT

Although the technology scaling has enabled designers to integrate a large number of processors onto a single chip realizing chip multi-processor (CMP), problems arising from technology scaling have made power reduction an important design issue. Since interconnection networks dissipate a significant portion of the total system power budget, it is desirable to consider interconnection network's power efficiency when designing CMP. In this paper, we present a variable frequency link for a power-aware interconnection network using the clock boosting mechanism, and apply a dynamic frequency scaling (DFS) policy, that judiciously adjusts link frequency based on link utilization parameter. Experimental result shows that history-based DFS successfully adjusts link frequency to track actual link utilization over time, demonstrating the feasibility of the proposed link as a power-aware interconnection network for system-on-chip (SoC).

1. Introduction

The technology scaling has enabled designers to integrate a large number of processors onto a single chip, realizing chip multi-processor (CMP). High performance CMP architectures have been gaining the attention of high performance computing community in the past few years. As the demand for network bandwidth increases for CMP, the idea of network-on-chip (NoC) becomes more promising because of performance, power, and scalability requirements for an SoC design [1].

Although today's processors are much faster and far more versatile than their predecessors using high-speed circuits and parallel processing, they also consume a lot of power. Moreover, an interconnection network dissipates a significant fraction of the total system power budget. For instance, the MIT raw on-chip network consumes 36% of the total chip power and Alpha 21364 microprocessor dissipates 20% of power in interconnection network [2]. Therefore, an interconnection network must be designed to be power-aware.

In this paper, we motivate the use of dynamic frequency scaling (DFS) link, where the frequency is dynamically adjusted to minimize power dissipation while maintaining the performance demands. First, we propose a novel DFS link which adopts a clock boosting mechanism [3], providing fast response time for frequency transitions and low hardware overhead. Next, a DFS policy is introduced that includes the link utilization estimator, DFS algorithm, and link controller. Finally, the DFS policy is applied to the proposed DFS link, demonstrating the power saving in an on-chip interconnection network. To the best of our knowledge, this is the first investigation of power reduction for on-chip interconnection network based on the clock boosting mechanism.

The rest of this paper is organized as follows. Section 2 addresses existing work on power saving for on-chip interconnection network. The proposed DFS link based on the clock boosting mechanism is introduced in Section 3. The implementation and the experimental results are presented in Sections 4 and 5, respectively. Finally, conclusions are drawn in Section 6.

2. Backgrounds

2.1. Dynamic voltage/frequency scaling

A communication link in NoC is capable of scaling power consumption gracefully commensurate with traffic workload. This scalability allows for the efficient execution of energy-agile algorithms. Suppose that a link can be clocked at any nominal rate up to certain maximum value. This implies that different levels of power will be consumed for different clock frequencies. One option would be to clock all the links at the same rate to meet the throughput requirements. However, if there was only one link in the design that required to be clocked at a high rate, the other links could be clocked at a lower rate, consuming less power.

The total power consumption in an SoC is the combination of dynamic and static sources. In this paper, our focus is on the

* Corresponding author. Tel.: +1 949 422 7130.
E-mail addresses: seunglee@uci.edu (S.E. Lee), nader@uci.edu (N. Bagherzadeh).

0167-9260/$ - see front matter © 2009 Elsevier B.V. All rights reserved.
doi:10.1016/j.vlsi.2009.01.002

dynamic power consumption which arises from circuit switching activity, due to charging and discharging of the switched capacitance. The dynamic power consumption depends on four parameters: a switching activity factor (α), physical capacitance (C), supply voltage (V), and the clock frequency (f)

$$P_D = \frac{1}{2}\alpha C V^2 f \quad (1)$$

$$f_{max} = \eta \frac{(V - V_{th})^{\beta}}{V} \quad (2)$$

Eq. (2) establishes the relationship between the supply voltage V and the maximum operating frequency f_{max}, where V_{th} is the threshold voltage, and η and β are experimentally derived constants.

Dynamic power consumption can be reduced by lowering the supply voltage. This requires reducing the clock frequency accordingly to compensate for the additional gate delay due to the lower voltage. The use of this approach in run-time, which is called dynamic voltage scaling (DVS), addresses the problem of how to adjust the supply voltage and clock frequency of the link according to the traffic level. The basic idea is that because of high variance in network traffic, when a link is under-utilized, the link can be slowed down without affecting performance. However, DVS requires thousands of clock cycles during transition between voltage levels and additional hardware overhead for each link.

The other way to manage power consumption is DFS. DFS only adapts the system clock frequency by setting all links in the network to the same voltage, but it does not always reduce the total energy consumption. For instance, the power consumed by a network can be reduced by halving the operating clock frequency, but if it takes as long to forward the same amount of data, the total energy consumed will be similar. DFS is valid when the target system does not support DVS or the goal is to reduce peak or average power dissipation, indirectly reducing the chip's temperature [4]. An alternative to save link power is to add hardware such that a link can be powered down when it is not used heavily.

2.2. Related works

System level power management has been applied to some interconnection networks. Wei and Kim proposed chip-to-chip parallel [5] and serial [6] link design techniques where links can operate at different voltage and frequency levels. When link frequency is adjusted, supply voltage can track to the lower suitable value. Although this link was not designed for both dynamic voltage and frequency settings, previously the link architecture was used for DVS link model.

There are three kinds of approaches for DVS. One is an on-line scheme which adjusts the link speed dynamically, based on a hardware prediction mechanism by observing past link traffic activities. Shang et al. [7] developed a history-based DVS policy which adjusts operating voltage and clock frequency of a link according to the utilization of link and input buffer. Worm et al. [8] proposed an adaptive low-power transmission scheme for on-chip networks. They minimized the energy required for reliable communication, while satisfying QoS constraints. One of the potential problems with hardware prediction scheme is that a misprediction of traffic can be costly from performance and power perspectives.

Li et al. [9] proposed a compiler-driven approach where a compiler analyzes application code and extracts communication patterns among parallel processors. These patterns and the inherent data dependency information of the underlying code help the compiler decide the optimal voltage/frequency to be used for communication links at a given time frame. Shin and Kim [10] proposed an off-line link speed assignment algorithm for energy-efficient NoC. Given the task graph of a periodic real-time application, the algorithm assigns an appropriate communication speed to each link, while guaranteeing the timing constraints of real-time applications.

Soteriou et al. [11] proposed a software-directed methodology that extends parallel compiler flow in order to construct a power-aware interconnection network, by combining both on-line and off-line approaches. However, current DVS techniques require not only thousands of clock cycles to shift between voltage levels, limiting their ability to respond to high frequency changes in network bandwidth demands [12], but also additional hardware overhead such as transmitter, receiver, PLL, and adaptive power supply regulator for each link.

Kim et al. [13] proposed dynamic link shutdown (DLS), which powers down links intelligently when their utilizations are below a certain threshold level and a subset of highly used links can provide connectivity in the network. An adaptive routing strategy that intelligently uses a subset of links for communication was proposed, thereby facilitating DLS for minimizing energy consumption. Soteriou and Peh [2] explored the design space for communication channel turn-on/off based on a dynamic power management technique depending on hardware counter measurement obtained from the network during run-time. Chen et al. [14] introduced a compiler-directed approach, which increases the idle periods of communication channels by reusing the same set of channels for as many communication messages as possible. Li et al. [15] proposed a compiler-directed technique in order to turn off the communication channels to reduce NoC energy consumption. Even though it saves power significantly during idle period, it has reactivation penalty including delay and additional power consumption during a transition.

Hsu [16] saved 30% of power consumption in the MPEG core by applying DFS power management mechanism using only three frequency levels (25, 50, and 100 MHz). However, DFS was applied to a core, not to interconnection network, in a tile-based NoC architecture. To the best of our knowledge, this paper is the first proposal which addresses DFS for interconnection network.

The novel contributions of our work are:

(1) a DFS link proposal for on-chip interconnection network which offers not only fast response time reducing the frequency transition penalty, but also reduces hardware cost, as compared to DVS link, suitable for system integration;
(2) use of narrow control period, as compared to conventional DVS control, reducing misprediction penalty that occurs in a hardware prediction scheme by adjusting the frequency more often;
(3) implementation of a variable frequency link that judiciously adjusts link frequency based on the link utilization estimation, reducing power consumption.

3. Variable frequency link

The clock boosting router was proposed to increase throughput and reduce latency of an adaptive wormhole router [3]. The key idea of clock boosting mechanism is *the use of different clocks in a head flit and body flits* because body flits can continue advancing along the reserved path that is already established by the head flit, while the head flit requires the support of complex logic, increasing critical path. Thus, it reduces latency and increases throughput of a router by applying faster clock frequency to a boosting clock in order to forward body flits.