

# Low Power Nanoscale Buffer Management for Network on Chip Routers

Suman K. Mandal	Ron Denton	Saraju P. Mohanty	Rabi N. Mahapatra
Texas A&M University	Texas A&M University	University of North Texas	Texas A&M University
College Station, TX	College Station, TX	Denton, TX	College Station, TX
skmandal@cse.tamu.edu	denton@cse.tamu.edu	Saraju.Mohanty@unt.edu	rabi@cse.tamu.edu

## ABSTRACT

Network-on-Chip (NoC) is an on-chip communication solution in the future system-on-a-chip (SoC) necessitating high performance operation with low power dissipation. We present a novel dynamic power management technique for low power NoC router buffers using nano CMOS SRAMS. A feedback controller was designed for block level power management and a power aware adaptive controller was designed for low power flit storage encoding to reduce energy consumptions in the router buffers. Experiments with the proposed scheme showed up to 20% reduction in energy consumption while improving throughput by up to 21%.

## Categories and Subject Descriptors

B.3.1 [Semiconductor Memories]: Static Memory (SRAM)

B.4.1 [Input/Output and Data Communications]: Communications Devices –Receivers, Transmitters

## General Terms

Algorithms, Management, Performance, Design.

## Keywords

Nanoscale Technology NoC, SoC, Router, Dynamic Power Management

## 1. INTRODUCTION

Network-on-Chip (NoC) has emerged as the favorite on-chip communication solution. Buffer size and allocation policy play an important role in the performance and efficiency of a NoC router [6][7]. Furthermore, studies have shown that buffers can consume as much as up to 79% of NoC router power [3]. Thus efficient management is necessary to ensure high performance and low power. Efficient schemes use SRAM arrays for their simplicity and high performance [5].

Nanoscale SRAM buffers are very suitable for NoC router design because of their speed, density and reliability. Power dissipation characteristics of Nanoscale SRAMs are unique. Traditional low power design techniques are not sufficient to ensure minimum power operation. To that end, a dynamic power management technique specifically designed for Nanoscale SRAM buffers is necessary. Such management technique will make use of buffer

allocation information and the Nanoscale SRAM power dissipation characteristics to minimize both static and dynamic power consumption while maintaining performance.

It is notable that the buffer utilization in NoC router is dependent on network congestion. Depending on the application communication pattern a given routers buffer utilization will vary over time. To provision for high utilization case it is necessary to provide enough buffer in each router. However, often the buffers are not utilized and remain idle and consume power. To avoid this, we propose dynamic block level buffer power management. To be able to benefit from this scheme it is necessary to use a central buffering strategy in the router. An example design is described in [5]. We propose a buffer design where the new flits are buffered in sets. Each set can hold some number of flits and is powered by a single source which can be turned on or off using a power gate. Hence depending on usage, the buffers can be turned on/off set by set. The number of active sets required to ensure zero performance hit is determined using a feedback controller.

Another observation about Nanoscale SRAM cells is that, storage of 0 and 1 are significantly different in terms of power consumption. This characteristic is exploited at per flit level to minimize power consumption during storage and read write. This is done by selectively inverting the flits based on their zero density. The decision to invert or not is taken using a simple adaptive controller.

The contributions of this paper are as follows:

- A *Feedback Controlled Block level Buffer Management is proposed for dynamic power management*
- An *Adaptive controller for efficient Flit level Power management is proposed*
- *Both power management techniques are thoroughly evaluated for performance and energy efficiency and showed to outperform static allocation by 21% increase in throughput and 20% reduction in energy consumption.*

The rest of the paper is organized as follows. Section 2 discusses related works in the literature. In Section 3, NoC traffic flow was characterized and the Nano-CMOS Buffer design was discussed. The Router design is presented in Section 4. The proposed dynamic power management techniques are presented in Section 5. Experimental evaluation of the design was presented in Section 6. Section 7 concludes the paper.

## 2. RELATED RESEARCH

Both circuit level and system level techniques have been proposed for NoC power management. There have been significant research works on router buffer power management for low power. Detail discussion on existing designs have been done in [1][3][6].

Zhang et al. have proposed a centralized buffer management to achieve enhanced buffer utilization [5]. Their scheme

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GLSVLSI'10, May 16–18, 2010, Providence, Rhode Island, USA.

Copyright 2010 ACM 978-1-4503-0012-4/10/05...\$10.00.

demonstrated a 50% decrease in total buffer requirement in their router. However, they did not provide an active power management strategy which can further reduce dynamic power. The proposed power management technique explores this possibility in central buffer router design to achieve superior power/performance characteristics.

Wang et al. have proposed a zero-efficient design for router buffers that optimizes the circuit level design of router buffer to minimize energy consumption [2]. The basis of their work has been the predominance of zeros in the NoC traffic. This is primarily a circuit level work under the assumption of high zero density and does not necessarily fare well when there is majority of one. Also they do not consider any system level information or active power management technique to adapt to the dynamic nature of the traffic. The proposed scheme differs from this in the way buffers are allocated dynamically and also the way flits are encoded while storage.

### 3. PRELIMINARIES

#### 3.1 NoC Flow Analysis

Synthetic traffic has been widely used to evaluate NoC architectures. However, an analysis of actual application communication patterns results in interesting observations that can facilitate advanced management techniques for low static and dynamic power operation. Using a NoC-based, full SoC simulator, a series of application benchmarks were executed and the traffic flow in the NoC was analyzed. The experiments and observations are discussed below.

##### 3.1.1 Experimental Setup

We use a full system simulation environment that models a flit accurate NoC to analyze the traffic flows while running real benchmarks. Further details on the simulation framework can be found in Section 6.1. We set up three sets of experiment to analyze the traffic pattern in a typical NoC based SoC.

###### 3.1.1.1 Topologies

A 16 node ring network was simulated with 8 processors, 4 memory modules and 4 producer consumers mimicking other devices.

2D Mesh is the most popular NoC topology in literature due to its simplicity and regular nature. The Mesh was also configured with same 8 processors, 4 memories and 4 producer consumer cores.

Similar to 2D mesh is 2D torus. Over 2D mesh, 2D torus has the benefit of a lower network diameter but suffers from increased link count. The core assignment is the same as in 2D mesh case.

###### 3.1.1.2 Applications

The applications were taken from the benchmark Suite Mediabench, Mibench and SPEC2006 [11][12][13]. Selected applications from these benchmarks were mapped on to the above mentioned topologies and experiments were performed.

##### 3.1.2 Flow Characterization

From the experimental results common NoC traffic flow was characterized as illustrated below. The results showed that majority of NoC utilization was seen in the burst mode transfers. Figure 1 shows the distribution of the different flow groups in terms of burst lengths. The benchmark traffic was dominated by the longer bursts [4 packets or more].




Figure 1: Distribution of Traffic Types for MediaBench applications

##### 3.1.3 Buffer Occupancy Analysis




Figure 2: Buffer Utilization Distribution in a Mesh Topology

Figure 2 shows the buffer occupancy analysis for the mesh topology described earlier. This result clearly demonstrates that even though the average utilization of the buffer space is low, all routers have experienced high peaks of buffer utilization. Due to this nature, a straight forward static buffer allocation based on the average utilization will cause performance degradation [See Section 6.2.2].

To mitigate this problem and optimize the buffer utilization a dynamic technique is necessary. This is the key motivation behind the proposed feedback controlled dynamic buffer management.


### 3.2 Nanoscale CMOS Buffer Design

In addition to efficient management technique, inherent efficiency in the buffer design itself plays an important role in the overall energy efficiency of the buffer. Ideally, the design of buffers has to be low power, fast and reliable. Between the alternatives of buffer design, such as SRAM and Registers, SRAM is advantageous for speed and power. A novel failure tolerant nanoscale CMOS SRAM based buffer design is adopted. The unique characteristics that make it particularly suitable for NoC router buffer design are analyzed in the rest of the section.

#### 3.2.1 Nanoscale CMOS SRAM Buffer

Single-ended SRAMs are known for their tremendous potential of low power dissipation. A seven transistor SRAM cell is shown in Figure 3(A) [18]. The SRAM cell is composed of a read and write access transistor (transistor 1), two inverters (transistors 2, 3, 4 and 5) connected back to back in a closed loop fashion in order to store the 1-bit information and a transmission gate (transistors 6 and 7). The transmission gate opens the feedback connection between inverters during the write operation. The cell operates on a single bit-line, instead of having two bit-lines as in standard six transistor SRAM cell. Both reading and writing operations are performed over the single bit-line. The word-line (WL) is asserted


high prior to write and read operation as similar to standard six transistor SRAM cell. When the cell is in a hold mode, the word-line is low and a strong feedback is provided to the cross-coupled inverters by the transmission.



**Figure 3: (A) Single-Ended 7-Transistor SRAM which is tolerant to failures induced by nanoscale process variations. (B) Current paths for write '0' operation of the 7T-SRAM.**

The total power dissipation of a CMOS based SRAM circuit for sub-65nm technology node is defined as the summation of dynamic power dissipation, subthreshold leakage, and gate-oxide leakage. The SRAM cells have a tendency to retain data for some duration of time as they cannot be shut off. The current flow (or power dissipation) in each device depends of the location the device in the SRAM circuit as well the operation (e.g. read, write, or hold) being performed. Thus, for accurate measurement of current (power) it is important that the currents are identified. Figure 3(B) shows the current paths for 'write 0' operation. The SRAM cell is simulated at the 45nm CMOS technology node using Predictive Technology Model [15] for nominal sized transistors at a supply voltage of 0.7V. The simulation results are presented in Figure 4 for the above 7-transistor SRAM when designed using dual-threshold voltage technique for low-power dissipation.

### 3.2.2 Statistical Power Model



**Figure 4: Total Power Dissipation of 7T SRAM**

In nanoscale CMOS process variations is a major concern. The process variation has made the designers job much complicated due to loss of circuit yield with reduced time to market. We have selected twelve process parameters for statistical variability study: NMOS/PMOS channel length, NMOS/PMOS channel doping concentration, access-transistor length and width, driver-transistor length and width, load-transistor length and width. Some of the parameters are independent and some are correlated which is taken into account during simulation for realistic study. Each of the process parameters is assumed to have a Gaussian distribution with mean ( $\mu$ ) taken as the nominal values specified in the PTM for 45nm node and standard deviation ( $\sigma$ ) as 5% of the mean. The

statistical process variation in parameters is translated to power, leakage, and Static Noise Margin using Monte Carlo simulations. Monte Carlo simulation is an efficient approach because it does not require relating the output to input which otherwise would have been cumbersome for the large number of parameters. For brevity, the statistical distributions of total power dissipation due to nanoscale process variations averaged over different operations are presented in Figure 4.

The static and dynamic power dissipation of the SRAM for different modes of operations is presented in Table 2. The probability density functions of all these are Gaussian in nature.

**Table 1: Static and Dynamic Power Dissipation of SRAM.**

Power	Operation	Mean ( $\mu$ )	Standard Deviation ( $\sigma$ )
Gate Leakage	Write 1	21.2nW	9.4nW
	Write 0	21.9nW	9.5nW
	Read 1	12.9nW	5.4nW
	Read 0	7.8nW	3.2nW
	Store 1	2.8nW	1.8nW
	Store 0	1.0nW	0.5nW
Subthreshold Leakage	Write 1	38.2nW	21.1nW
	Write 0	7.8nW	19.0nW
	Read 1	12.3nW	27.0nW
	Read 0	13.5nW	32.1nW
	Store 1	10.8nW	21.0nW
	Store 0	16.2nW	2.3nW
Dynamic Power	Write 1	39.2nW	22.1nW
	Write 0	5.1nW	20.0nW
	Read 1	14.3nW	30.0nW
	Read 0	15.5nW	32.1nW
	Store 1	12.8nW	22.0nW
	Store 0	17.2nW	2.9nW

## 4. ROUTER ARCHITECTURE


The proposed buffer design is very suitable for routers with centralized buffer management. In this section we discuss the design of a single cycle centralized buffer router design.

### 4.1 Virtual Buffer Architecture

Figure 5 illustrates the Router design. To effectively utilize the central buffer design a concept of virtual buffer is introduced. Every input port contains virtual buffer in which each valid entry points to a queue in the central physical buffer. Queue management is performed in the physical buffer design. Instead of allocating the buffer based on virtual channels, a concept of set and line is introduced. A set is a collection of lines allocated to packets going to a given output port of the router. So Set 0 in any input port will contain packets that intend to go to Output Port 0. This requires one step look-ahead routing. Each line queues packets that are going to the same destination. This property avoids head of the line blocking.

### 4.2 Central Physical Buffer Design

The virtual buffers allow independent management of the central buffer structure. The physical buffer is managed centrally and each virtual buffer may or may not be mapped to a physical buffer. To be able to effectively perform power management using power gating the buffer is grouped in blocks. Each block can be turned on and off using a power gating structure. This is desired compared to turning on/off each buffer element because of the power gating structure overhead. The available buffer index logic selects free buffer elements from the fullest buffer block which is not full. This leads to buffer blocks being utilized one by one. Hence allowing blocks to be turned off when no buffer



**Figure 5: The Central Buffer Router Architecture**


element from that block is used. This can be implemented using a priority encoder based combinational logic block or a PLA.

## 5. BUFFER POWER MANAGEMENT

The dynamic power management is motivated by the traffic flow analysis presented in Section 3.1.2. It was shown that bursts account for a large proportion of network traffic and also the bursts are in general restricted to a few network paths.

Taking advantage of these characteristics, a mixed mode feedback controller was designed to do buffer power management at block level. To further enhance the control on power consumption an adaptive controller for flit storage encoding management was introduced. We will discuss each level of power management in the following sections.

### 5.1 Block Level Power Management




**Figure 6: Block Level Feedback System**

A non linear feedback controller was designed for block level power management. Figure 6 shows the feedback system modeling. The observed traffic ( $\lambda'$ ) is represented by a function of the injection load ( $\lambda$ ) and the backpressure ( $f$ ) of the network which is in turn again a function of the available buffer space. The feedback function  $f()$  is estimated from simulation and tabulated. This definition is used to calculate the minimum buffers required to maintain performance.

#### 5.1.1 The Feedback Function

To predict the buffer utilization, a flow density metric is used. The flow density of a given set represents the likelihood of that set being occupied. The flow density of a set is given by number of flows \* bandwidth of each flow / available buffer for that set.

#### 5.1.2 Controller FSM



**Figure 7: The Block Level Power Manager FSM**

The block level power management is done by utilizing flow prediction and buffer power gating. Figure 7 shows the simple FSM used to do this operation. The timeout and the threshold of update are set based on the dynamic nature of the application setup. System level support is used here in predicting flow densities. Every packet is marked as Start Burst, End Burst or Random. Incoming Start/End of Burst packets increase the confidence of prediction based on the flow rate and the length of the burst, while Random packets reduce the prediction confidence. This is utilized to dynamically adjust the buffer switch threshold. The buffer resize is done in a slow mode not to cause power surges. The shown FSM sets a register with the new required buffer number. The power controller shuts down buffers one by one until the remaining buffer matches the required number. Same procedure is followed while increasing the buffers. Buffer blocks are turned on one by one.

## 5.2 Flit Level Power Management

In addition to the block level buffer management, a dynamic encoding technique is applied per flit to further enhance the energy efficiency. This is done by utilizing either positive or negative logic based on the zero density of the data in the flits.

### 5.2.1 Flit Storage States

Any flit can be stored in one of the three states: Active 0, Active 1 or Sleep. In Active 0 – true logic is stored as 0. In Active 1 – true logic is stored as 1. In sleep state data is not stored. Hence sleep is a non-preserving state. A linear adaptive control mechanism is designed to assign the flit storage states dynamically. Wrapper logic is added in the buffer design to make this process transparent to the rest of the system.

### 5.2.2 Adaptive Control

An adaptive control technique was developed for the flit level power management. The overall control operation is shown in Figure 8. The state is a simplified representation of the density of 1's in the flit.




Figure 8: Adaptive Control for State Assignment

### 5.2.3 Estimator Design

For low overhead the estimator needs to be simple. In the proposed design the flits are marked to be ‘1-dense’ by adding a bit to the header. This bit is set when the flit is created. A simple estimate is the frequency of this bit being set in a given time interval  $T$ . The corresponding estimator can be easily implemented using a saturation counter.

### 5.2.4 Controller FSM

The estimator described above makes the controller design very simple. The estimate is updated every time a new flit comes in. After every time interval  $T$ , this estimate is compared with a threshold. If the estimate is higher than the threshold the flit is inverted when stored. This decision remains for the next  $T$  time. After which the condition is re-evaluated. Figure 9 depicts the operation in a flowchart.




Figure 9: FSM for Dynamic Flit Inversion Controller

## 6. EXPERIMENTAL RESULTS

The proposed dynamic buffer management technique was compared with a static buffer allocation. The evaluation consists

of experiments to demonstrate performance in terms of latency and throughput, power efficiency and design overhead in terms of area and power. The simulation platform and the experimental results are discussed in the following section.

### 6.1 Simulation Platform

We used NoCSim [16], a flit accurate NoC simulator capable of running SparcV8 processors with integrated cache and memory model to simulate the proposed router design. The simulator is written in SystemC for flexible modeling detail and simulation speed. To simulate real benchmark programs, the ArchC Instruction Set Simulator [15] was integrated with the NoC simulation framework. This simulation system allows of full SoC simulation with a software system kernel. A 9 core mesh system with 5 processors at center and 4 memory cores in four corners was used to evaluate the proposed buffer management scheme.

### 6.2 Performance

#### 6.2.1 Latency




Figure 10: Latency vs. Injection Rate with The three schemes

To measure the effect of the proposed buffer management schemes experiments were performed using statically allocated buffers based on average utilization, statically allocated buffer based on maximum buffer utilization and dynamically managed buffer with the proposed scheme. Figure 10 compares the three schemes based on end-to-end latency. The dynamically managed buffer allocation results in latency comparable to the average or max allocation case. Note that maximum buffer allocation increases the latency toward the higher injection rate. This happens due to increased contention in the router because more flits are buffered.

#### 6.2.2 Throughput





Figure 11: Throughput Comparison

Figure 11 compares the throughput achieved by the three schemes based on varied injection rate. The dynamic buffer allocation

achieves virtually the same throughput as the maximum buffer allocation case across the range of injection rate. But the static allocation based on average buffer utilization takes major hit in terms of throughput beyond 30% flit injection rate. And it deteriorates with higher injection rate. At saturation (0.56) the throughput reduction is as much as 21%.

### 6.3 Power



**Figure 12: Energy saving using Dynamic Allocation and Adaptive Flit Inversion over Static Allocation**

Figure 12 illustrates the energy savings achieved by using the dynamic feedback controlled buffer allocation and the adaptive flit storage encoding. The combined technique achieved up to 20% energy saving compared to static buffer allocation technique at 30% injection rate. At higher flit injection rates congestion causes frequent change in the buffer requirement thus leading to repeated adjustment of buffer allocation resulting in slightly higher energy consumption. Also, frequent on/off of the buffers causes more write overhead in the encoding scheme and hence reduces energy savings little further.

### 6.4 Overhead

The proposed feedback controller for the block level buffer management and the adaptive controller for flit level power management was designed and synthesized using 90nm TSMC Library to calculate the area and power overheads shown below and are minimal. 90nm technology was used due to unavailability of sufficient library support for 45nm. However, the results are indicative of the low overall area & power overhead.

**Table 2: Controller Overheads**

Overhead	Buffer Allocator	Flit Encoding Selector	Total
Area	870 GE	1172 GE	2042 GE
Power	91 $\mu$ W	69 $\mu$ W	160 $\mu$ W

## 7. CONCLUSION

A novel low power nano CMOS buffer design was presented. The proposed technique utilizes system level information to do effective dynamic power management of the router buffer. Experimental evaluation have demonstrated the proposed design to be outperforming static buffer allocation by 21% in terms of throughput while consuming up to 20% less power.

## 8. REFERENCES

[1] Simunic, T., Boyd, S., "Managing power consumption in networks on chips," Design, Automation and Test in Europe Conference and Exhibition, pp.110-116, 2002.  
 [2] Jun Wang; Hongbo Zeng; Kun Huang; Ge Zhang; Yan Tang, "Zero-Efficient Buffer Design for Reliable Network-on-Chip in Tiled Chip-

Multi-Processor," Design, Automation and Test in Europe, pp.792-795, 2008.  
 [3] Banerjee, A., Mullins, R., and Moore, S. "A Power and Energy Exploration of Network-on-Chip Architectures", In Proceedings of the First international Symposium on Networks-on-Chip, pp. 163-172, 2007.  
 [4] Young Hoon Kang; Taek-Jun Kwon; Draper, J., "Dynamic packet fragmentation for increased virtual channel utilization in on-chip routers," ACM/IEEE International Symposium on Networks-on-Chip, pp.250-255, 2009.  
 [5] Wang, L., Zhang, J., Yang, X., and Wen, D., "Router with centralized buffer for network-on-chip". In Proceedings of the 19th ACM Great Lakes Symposium on VLSI, pp. 469-474, 2009.  
 [6] Ogras, U. Y., Hu, J., and Marculescu, R., "Key research problems in NoC design: a holistic perspective.", In Proceedings of the 3rd IEEE/ACM/IFIP international Conference on Hardware/Software Codesign and System Synthesis, pp. 69-74, 2005.  
 [7] Chrysostomos A. Nicopoulos, Dongkook Park, Jongman Kim, N. Vijaykrishnan, Mazin S. Yousif, Chita R. Das, "ViChaR: A Dynamic Virtual Channel Regulator for Network-on-Chip Routers," *IEEE/ACM International Symposium on Microarchitecture*, pp. 333-346, 2006.  
 [8] Shim, K. S., Cho, M. H., Kinsy, M., Wen, T., Lis, M., Suh, G. E., and Devadas, S. 2009. Static virtual channel allocation in oblivious routing. In *Proceedings of the 2009 3rd ACM/IEEE international Symposium on Networks-on-Chip*, pp. 38-43, 2009.  
 [9] Kodi, A., Sarathy, A., and Louri, A., Design of adaptive communication channel buffers for low-power area-efficient network-on-chip architecture. In *Proceedings of the 3rd ACM/IEEE Symposium on Architecture For Networking and Communications Systems*, pp. 47-56, 2007.  
 [10] Koranne, S., "Element Interconnect Bus", Practical Computing on the Cell Broadband Engine, Springer US, pp. 61-66, 2009 (ISBN: 978-1-4419-0342-6).  
 [11] Lee, C., Potkonjak, M., and Mangione-Smith, W. H. 1997. MediaBench: a tool for evaluating and synthesizing multimedia and communications systems. In *Proceedings of the 30th Annual ACM/IEEE international Symposium on Microarchitecture*, pp. 330-335, 1997.  
 [12] M.R. Guthaus, J.S. Ringenberg, D. Ernst, T.M. Austin, T. Mudge, R.B. Brown, "MiBench: A free, commercially representative embedded benchmark suite," *Annual IEEE International Workshop Workload Characterization*, pp. 3-14, 2001.  
 [13] The SPEC2006 Benchmarks, <http://www.spec.org>.  
 [14] Kodi, A., Sarathy, A., and Louri, A., Design of adaptive communication channel buffers for low-power area-efficient network-on-chip architecture. In *Proceedings of the 3rd ACM/IEEE Symposium on Architecture For Networking and Communications Systems*, pp. 47-56, 2007.  
 [15] W. Zhao and Y. Cao. New Generation of Predictive Technology Model for sub-45nm Design Exploration. In Proceedings of International Symposium on Quality Electronic Design, pp. 585–590, 2006.  
 [16] <http://archc.sourceforge.net>.  
 [17] S. K. Mandal, N. Gupta, A. Mandal, J. Malave, J. D. Lee and R. Mahapatra, "NoCBench: A Benchmarking Platform for Network on Chip", In Proc. of Workshop on Unique Chips and Systems, 2009.  
 [18] G. Thakral, S. P. Mohanty, D. Ghai, and D. K. Pradhan, "A Combined DOE-ILP Based Power and Read Stability Optimization in Nano-CMOS SRAM", in *Proceedings of the 23rd IEEE International Conference on VLSI Design*, pp. 45-50, 2010.